STABILITY MODELING AND GAME-THEORETIC CONSIDERATIONS

Rick Picard, Todd Graves, Jane Booker

1. Stability Measures

The potential drawdown of U.S. and Russian nuclear forces raises issues of future "stability." Once these two powers no longer have arsenals that are as dominant when compared with the rest of the world as they once were, how will that affect global stability? Armchair arguments exist in both directions, such as

a) The U.S. and Russia would be less inclined to use nuclear weapons on each other because expending their reduced arsenals could leave them vulnerable to attacks from other powers (e.g., China). Hence a stable situation follows if the drawdown were handled properly; or

b) Coalitions could arise that would become threats because a smaller combined nuclear arsenal would be required for them to do so, possibly leading to an unstable situation. Moreover, certain countries (Germany, Japan, Canada) might develop their own nuclear arsenals if they felt that reduced U.S. nuclear forces were no longer adequate to protect them, while others (Israel, India, Pakistan) might not feel as inhibited in using nuclear weapons to pursue their individual future agendas.

To assess matters, a formal definition of stability, together with quantitative metrics to assess them, would be helpful.

Game theory provides one basis for such analysis. The subject has been extensively studied relative to arms control issues for some time (see bibliography). Though game theoretic explanations of human/organizational behavior are not perfect, they allow for a formal exposition of the underlying thought processes, value systems, and so on, behind any claimed conclusions. As such, a better understanding of the situation can evolve.

Relative to so-called "crisis stability" and "geopolitical stability," game theory has serious limitations that are discussed in the sections to follow. Basically, most realistic arms control problems cannot be so neatly pigeonholed as required by game theory, with each side having nearly complete knowledge of the other's motivations and available actions. This is why, despite the vast literature on the subject, game theory is not recognized as *the* answer to understanding behavior. In other domains, such as “force stability,” the action spaces and payoff matrices have stronger underpinnings, and stability can be examined in some detail.

In this report, we briefly review the existing game theoretic literature, discuss the strengths/weaknesses of applying game theory to certain arms control environments, and illustrate application with a simple example in force stability. Treatment of the stochastic aspects of the example goes beyond that found in the current literature for game theory and illustrates how the usual concepts (minimax solutions, equilibrium points) must be extended in a stochastic environment.

The presentation is self contained, and makes no assumption that the reader is already familiar with the subject matter.

2. Benefits of Game Theoretic Modeling of Behavior

Game theory provides a logical basis for decision making and for defining different types of stability. One of its earliest advocates, interestingly enough, was prominent for his work at Los Alamos — John Von Neumann. This subject has been studied a great length, and is a useful means of organizing the considerations relevant to decision problems. A vast literature exists on the methodology (see the bibliography for a sampling of related work), and we do not attempt a complete summary. Instead, basic concepts are discussed relative to arms control issues and, specifically, in their relation to stability.

To introduce the subject at an introductory level, consider a game theoretic model of the children’s game "rock-paper-scissors." The payoff matrix for the game is listed below, where the first entry in each cell is the payoff to Player 1 and the second entry in each cell is the payoff to Player 2. Each player can choose from among the three actions (Rock, Paper, or Scissors), and the players’ actions are announced simultaneously. For example, if Player 1 chooses Rock and Player 2 chooses Scissors then Player 1 wins (his payoff is +1) and Player 2 loses (his payoff is -1). If both players choose the same action, neither wins (both payoffs are 0). The payoffs for all possible combinations of actions are listed in the table.

		Player 2's Action:		
		Rock	Paper	Scissors
Player 1's Action:	Rock	0, 0	-1, +1	+1, -1
	Paper	+1, -1	0, 0	-1, +1
	Scissors	-1, +1	+1, -1	0, 0

If one player can outguess the other, he wins the game. But without knowledge of the other player's tendencies, there is no formal solution to the rock-paper-scissors game. That is, there is no single best choice of action that offers either player an advantage.

Indeed, if one player were to adopt the randomized strategy of choosing each possible action with probability 1/3, it no longer matters what the other player does: the expected payoff is 0 for both players. In the context of a sequential game, where the game is repeated multiple times with the same payoff matrix, independent randomization for each game offers a guaranteed payoff, at least in a statistical sense, and provides one type of a stable solution. Namely, that there is a stable, long term payoff against all opponents' strategies.

The rock-paper-scissors game also illustrates the problem with infinite regress, and a kind of reverse psychology. That is, suppose Player 1 "knows" that he plans to choose "Paper." If Player 2 knows that Player 1 knows he's planning to choose "Paper," then Player 2 will choose "Scissors." But if Player 1 knows that Player 2 knows that Player 1 knows he's planning to choose "Paper," then Player 1 will counter Player 2's anticipated "Scissors" action by choosing "Rock." However, if Player 2 knows that Player 1 knows that Player 2 knows that Player 1 knows . . . . This is a classic illustration of infinite regress, where each player can rationalize, ad infinitum, about the best strategy against an opponent having some intelligence regarding his planned activities. The infinite regress situation is unstable in some sense.

Consider a game where another type of stable solution exists. The payoff matrix for this game is:

		Player 2's Action: Low	Medium	High
Player 1's Action:	Low	+2, -2	-1, +1	+3, -3
	Medium	+1, -1	0, 0	+1, -1
	High	-3, +3	-1, +1	-2, +2

The combination of actions where each player chooses "Medium" is known as a Nash equilibrium. In distinct contrast to the rock-paper-scissors game, knowledge that the opponent will choose "Medium" offers neither player the opportunity to gain by changing his own action away from "Medium." The Nash equilibrium reflects a region of the payoff space where neither side is motivated to unilaterally change its behavior, and thus the equilibrium point constitutes one notion of stability.

The two payoff matrices above are examples of zero-sum games, where the sum of payoffs to Player 1 and Player 2 always add to zero. That is, each player can gain only at the expense of the other. In certain bi-polar settings, such as an oversimplified view of the cold war, this type of payoff matrix may be plausible. In most environments, however, the parties involved have mutual interests and the zero sum phenomenon is not realistic.

A third game, frequently discussed in the context of arms control, is the classic prisoner's dilemma. This is not a zero sum game. Background for one version of the prisoner's dilemma is as follows. The two players represent apprehended criminals, imprisoned in separate jail cells following the commission of a crime. The district attorney, knowing the evidence against them is too weak to obtain a stiff sentence, approaches the prisoners individually to offer each a plea bargain. Possible outcomes are:

- If one prisoner Confesses and the other Stonewalls, the confessor will receive probation without jail time, his evidence being used to obtain a stiff sentence for the stonewaller.
- If both prisoners Stonewall, both will receive light jail sentences because of the limited evidence against them.

- If both Confess, both will get medium sentences (receiving some leniency for sparing the district attorney the ordeal of a taking the case to trial).

An example of a payoff matrix for the prisoner's dilemma game is as follows:

		Player 2's Action:	
		Confess	Stonewall
Player 1's Action	Confess	-10, -10	0, -20
	Stonewall	-20, 0	-3, -3

One solution to the prisoner's dilemma involves a Nash equilibrium. Player 1, thinking selfishly, knows that no matter which action is taken by the other player, it's to his benefit to confess. Player 2, also reasoning selfishly, confesses as well, leaving both players with a payoff of -10. The solution is stable in the Nash sense, in that neither player, knowing that the other will confess, can gain by unilaterally changing his own action.

Although the "Confess-Confess" solution is optimal in the selfish sense described above, it leads to a foolish result collectively. *Both players* would be better off with the "Stonewall-Stonewall" solution than with the (supposedly) selfishly optimal "Confess-Confess" solution, if only they could cooperatively achieve it. If they could somehow enter into a binding agreement to Stonewall, each player could achieve a good resolution, without fear that the other player would exploit that action for personal gain.

In the games above, there is no (formal) stochastic component. Other than the consideration of so-called mixed strategies, such as randomly choosing each of rock-paper-scissors with probability 1/3 in order to guarantee a minimum statistical payoff, other aspects of the game are deterministic. There are certain other types of games having stochastic components, a few of which are reviewed here.

Perhaps the simplest such game is the "noisy duel." That is, two combatants, each having a gun containing a single bullet, square off in a duel. Initially, the combatants are distance D apart, and they slowly walk towards each other. When combatant i (i = 1, 2) fires his weapon, his probability of killing his opponent is $p_i(d)$, where d denotes the distance from his opponent at the time his gun is fired. If one combatant fires and misses, however, his opponent is certain to walk up to him and kill him.

The noisy duel game, then, reduces to each combatant deciding when to shoot. If someone shoots too soon, and with too low a probability, he may miss and lose the game. On the other hand, if he waits too long, his opponent may kill him first. In terms of striking the right balance, the optimal solution depends on the respective payoffs of the players (aside: for payoffs of ±1, the solution is for each player to fire as soon as $p_1(d) + p_2(d) = 1$).

Still other games depend on somewhat unpredictable psychological factors. Consider a one-dimensional game of hide-and-seek. The first player picks an integer between 1 and N, and the second player is to guess repeatedly until he correctly identifies that integer. Upon each incorrect guess, the second player is told whether his most recent guess is high, low, or correct. Using that information, he continues guessing until he guesses correctly. The first player would like to maximize the number of guesses needed, while the second would like to minimize the number.

An extension of this game is to introduce additional information — the players may know each other and believe that they can anticipate (however imperfectly) the other's action. That information can then be used in determining a sequential (and possibly mixed) strategy.

Still other games exist, involving more than 2 players, for example. As noted above, the purpose of this report is not to exhaustively review the considerable literature on the subject. Instead, it is to convey basic ideas from game theory to provide a normative theory for decision making and different notions of stability. Because all available actions for all players are detailed, the assumptions regarding the resulting outcomes are quantitatively described and the basis for selecting a strategy is open to examination by others. The structure supplied allows for the game to be objectively examined for optimal strategies and for evaluating various notions of stability. This is the greatest strength of the approach.

3. Recognized Limitations of Game Theoretic Modeling of Arms Control Behavior

The above (oversimplified) examples of two-person games were introduced for purposes of illustrating some basic points and different notions of stability. Note that these examples involved several ground rules which warrant further discussion. One perspective on these shortcomings is that they identify areas where game theoretic research would be useful in aiding study of arms control applications. Given the concerted effort to apply game theory to international stability (see the bibliography for a sampling of the literature), if current methods were adequate to address the situation, the problem would have been solved by now.

Game theoretic shortcomings include the following.

3.1 Action spaces are assumed known to all players.

Specification of the payoff matrix defining the game involves a complete enumeration of all possible actions that all players are allowed to take. In many realistic decision problems, options are not so well delineated. An action can involve many distinct sub-actions, which in turn involve sub-sub-actions, and so on. At some level of detail, a complete characterization of the action space is impossible, and only low-resolution approximations can be considered. Ideally, the summary incorporates all relevant considerations and is reasonable to first order, though there is rarely a way to verify this assumption.

History is littered with examples where imperfect military intelligence led to instances where countries have miscalculated the responses of other countries to their actions. Historical examples exist where adversaries have overestimated and underestimated each others' capabilities, and were surprised to see that their opponents were capable of taking the actions they did.

People act on a combination of beliefs, some of which may be correct and others not. In a related issue, information on which they base actions can be timely or outdated. Because of this aspect of uncertainty, decisionmaking is less like a well defined two-person game having a deterministic payoff matrix and more like a game of poker, where bluffing is an accepted (and sometimes successful) strategy. Although simple types of bluffing can be incorporated into action spaces, the more general case cannot be.

3.2 The payoffs to each player of all possible combinations of actions are deterministic and are known to all players.

3.2.1 Amalgamation and Stochastic Payoffs

Except in isolated dictatorships, governmental decisions are not truly made by a single individual. Policy decisions must take into account numerous considerations (e.g., political, economic, military, personal agendas) and differences of opinions may exist regarding the tradeoffs involved. Thus, behavior is affected when countries don't behave the way that individual people do. Just look at U.S. government purchasing rules, for example, in contrast to the way that individuals make personal buying decisions. Another example is that smaller organizations sometimes optimize their own interests at the cost to their larger organization. In the extreme, horses are designed by committee.

The multiple considerations are rolled into a single payoff value in order to force fit the problem into a game theoretic setting. In so doing, measurement units may be difficult to assess and then reconcile, such as a nation evaluating the tradeoffs between economic benefits to its society vis-a-vis devoting substantial national resources to maintaining a vast military machine. In extreme cases, issues become very problematic (e.g., how many dollars is a human life worth?). There is a literature on multi-attribute utility analysis (e.g., Keeney and Raiffa 1993; Krakowski 1996) which attempts to address this problem by assigning a "grand utility" to each outcome. But summarizing succinctly, "avoiding weasel wording, it boils down to the unhappy conclusion that there is no universal satisfactory way to amalgamate the individual preferences into group preferences" (Shubik, p. 100).

The issue of amalgamation aside, many decisionmaking situations involve payoffs that are not known with certainty. In war games and military simulation models, for example, stochastic elements are intrinsic to the game. If a decision is made to attack, for example, the outcome is not known with certainty. Moreover, subjective perceptions can be difficult to quantify, especially in regard to assigning risk to low-probability, high-consequence events (such as nuclear exchanges) and issues with high emotional content (for example, nuclear power), and there is some evidence (Fischoff, Slovic, and Lichtenstein 1982) to suggest that presumed experts in an area are often no better than the general public in making such assessments. Cultural differences are also important in this regard (Whitman 1985). As such, different people can react to the same situation differently and the subject of risk perception becomes important.

At a minimum, efforts to determine the sensitivity of the optimal strategy to perturbations of the payoff matrix are essential parts of any analysis.

Not only is assessing the payoffs to one's own side ill-defined, but assessing the payoffs to others introduces still more complications. Imperfect intelligence may exist regarding options available to the other side(s), as well as the effectiveness of those options if implemented. Indeed, one aspect of the counterintelligence involves feeding misinformation to adversaries. Moreover, the motivations of adversaries, as well as the value systems that drive their motivations, are frequently difficult to determine from a distance. The Cuban missile crisis is a classical case in point, where American decisionmakers spent considerable time (e.g., Allison 1971) trying to understand how the situation was being perceived by the other side. A subjectivist approach to the problem, formalizing subjective uncertainty through the use of probability distributions, leads to stochastic payoffs in much the same way as does modeling uncertainty in military attacks.

When the payoff matrix is stochastic, the usual game-theoretic notions of stability, such as minimax solutions and Nash equilibria, no longer exist in their pure form. For example, with certain stochastic models such as the normal distribution, there is no guaranteed minimum payoff. Minimums only exist in a probabilistic sense (e.g., "with probability 90%, the payoff is at least x."). At this point, utility functions can become important, and the notion of risk aversion affects the definition of "optimal" strategies. Use of expected values (Booker and Bryson 1985) is sometimes reasonable, and sometimes not. Possibly as a result of these factors, there is little literature on the subject of stochastic payoffs.

3.2.2 Knowedge Management

Much knowledge, or "know how" or expert insight, is not (fully) quantitative but still important for understanding and evaluating stability. In Section 3.2, the lack of knowledge by the players involved was noted as a shortcoming of the game theoretic approach. With modern day computation and communication capabilities, knowledge and other forms of factual information, including observational and collected data, can be gathered, documented, stored and retrieved for analysis with greater, ease.

Knowledge management is a burgeoning technology (see Knowledge Transfer International 1999). Its methods support the process of organizing, transferring, and using the heterogeneous information and expertise of groups of people engaged in an activity (such as operating a military force). Knowledge Management is commonly used to preserve and leverage knowledge, efficiently link seekers to information so that they can make informed decisions, enable diverse or dispersed groups to work towards a common purpose. Knowledge Management includes the elicitation and documentation of the knowledge of experts (knowledge acquisition), the creation of electronic repositories, such as knowledge bases or organizational memories, and use of tools for electronic collaboration, data searching, warehousing, mining, and discovery.

Perhaps with better knowledge management, game theory and these alternative theories can implemented to overcome the weakness of game theory in addressing certain decision problems.

3.3 Players behave rationally relative to their respective payoffs.

Some aspects of defining rational behavior have been discussed above. Issues of individual (selfish) rationality as opposed to collective rationality, such as was discussed with respect to the prisoner's dilemma, or to amalgamating many distinct points of view in a single payoff matrix, make it difficult to examine rationality. But even at the level of an individual person, rational behavior is difficult to address.

Unless rationality is circularly defined (i.e., whatever action was taken *must* have optimized the relevant utility function at the time), people frequently don't behave rationally. Either that, or they behave logically with respect to a value system that is so poorly understood (by their adversaries if not by themselves) that game theoretic modeling is highly problematic.

Irrational behavior may follow from (Meyer and Booker, 1991)

- Failure to update their beliefs in light of new information that becomes available to them – e.g., beliefs can change slowly in spite of large amounts of contrary evidence.
- Acting on perceptions that are in fact false.
- Acting on personal agendas.
- Poor estimation of frequencies of events, especially rare ones.
- Poor understanding of interdependencies between events.
- Underestimation of uncertainty – e.g., placing excessive confidence in their own beliefs or in the beliefs of (supposed) experts.
- Just plain stupidity.

Egos can also play a role. Shubik (1975), for example, presented a study where people turned down money only because they felt slighted by the process in which it was offered. Substantial person-to-person variability was also present in that study, illustrating that it is dangerous to generalize how one person will react given information on how others reacted under the same circumstances.

Attempts have been made (e.g., Zagare and Kilgour 1995, Giles, et. al. 1994) to assign psychological labels to different types of players and view payoffs for each in that light. This allows for characterizing different degrees of risk aversion versus risk taking, different pain thresholds (e.g., the willingness to absorb damage in order to inflict damage, important in military exchange simulations), and so on. Such a view provides for flexibility in capturing individualistic tendencies, although this flexibility comes at the cost of a well defined game when the tendencies of an adversary are unknown.

An alternative view, as stated by Weber (1991, p. 68), is that "Decisionmakers here do not deal in utility functions. Instead, they deal in arguments." By implication, force fitting decision problems into a game theoretic framework is ill advised.

3.4 Players have sufficient time to search the solution space and arrive at an optimal solution.

An aspect of decisionmaking that is poorly captured by game theory involves the cost of searching the solution space for an optimal result. As a simple example, consider a person making a purchasing decision. The person goes to a store, examines relevant items, and has the choice to buy an item that's available at the price given, or to continue the search at another store in the hope of obtaining a more favorable deal. If the search continues, the same buy-or-continue-to-look decision is again faced at the next store.

In this example, there is an implicit cost-benefit tradeoff of continuing the search for a better deal than has been available thus far (which involves continued investment of effort on behalf of the person) versus stopping the search and buying at a particular point in time. For major purchasing decisions (such as a house or automobile), a continued search may be worthwhile, whereas for minor purchasing decisions (such as a grocery item), it may be more cost effective to purchase immediately even if it is likely that a better deal could be found elsewhere. An assessment of the unknown — i.e., does a better deal than observed thus far even exist? — becomes a subjective element of the decision.

For international events, time pressures often exist. The old adage "he who hesitates is lost" may apply, but the tradeoffs in the continued pondering of the action space are less well defined. Actions to "buy time" may be taken, but in many cases (e.g., the Cuban missile crisis) it is not possible to think through the decision for as long as may be desired. At some point, it may be necessary to knowingly proceed with a suboptimal solution rather than spending additional time making minor improvements to it.

The effects of time pressure on decisions and the subsequent stability of a situation are potentially important but poorly understood. As has been noted (e.g., Shubik 1975) it is not realistic to conduct controlled experiments on this subject.

3.5 One-Time Payoffs.

In international relations, the “players” (nations, if not their specific leaders) involved have a long history of interacting with each other. The “game” is less of a one-time event, and more of an ongoing series of events with a corresponding series of payoffs. Circumstances and payoffs change over time.

To be sure, a literature exists on sequential games, usually assuming a constant payoff matrix. This was mentioned in the spirit of playing the rock-paper-scissors game multiple times. Another example is the so-called differential game (Isaacs 1965), which involves “lengthy sequences ... of decisions which are knit together logically so that a perceptible and calculable pattern prevails throughout.” Often, the conditions facing the players evolve according to a differential or difference equation, which gives the class of games its name. Examples of differential games are games of pursuit, in which one player chases another player across some domain, while the latter player attempts to avoid capture for as long as possible. The players receive periodic information as to the whereabouts of each other, and the payoff is dependent on the number of steps required for the capture to occur (in many games, capture is certain if the game is played long enough).

In many arms control environments, a complication is that the payoff matrix changes over time. As an example, the ABM treaty succeeded in the 1970s because the U.S. and (then) Soviet Union felt it was in their mutual interest. With the future prospects of other powers such as China and North Korea being able to strike with similar missiles, people are now rethinking the matter. As technology evolves and is redistributed across the world, circumstances change.

Such evolving situations can be modeled by examining multiple payoff matrices (i.e., a matrix for each time point of interest) or by assigning integrated (over time) payoffs for a single game. Another alternative is to use tree structures (e.g., Zagare and Kilgour 1995), but the efficacy of such structures is limited. In practice, the longer the scenario, the harder it is to anticipate an adversary.

Similar to a chess game, it becomes progressively more difficult to anticipate an opponent's actions 1, 2, 3, . . . moves into the future. Many factors affecting the rate of change in the action space are not directly observable (e.g., if a nation decides to develop nuclear weapons to allow it to have more flexibility in its actions, how much time is required for it to do so?). And the fact that many decisions involve multiple players (e.g., how do one nation's actions evolve from changes in another nation's actions?) only makes the time-dependent modeling of payoffs/actions more problematic.

3.6 Alternatives to Standard Quantitative Payoffs

Because of the shortcomings cited above in using standard game theory and quantitative payoffs, alternatives have been pursued in the literature that offer potential for capturing the human-based issues involved. Zadeh (1965) provides an alternative view of uncertainty by defining the concept of fuzzy logic, based on fuzzy set theory. In conventional logic and set theory (called crisp), each member of a population is assigned exactly to a set. For example, an F-16 is a military jet. In fuzzy logic, however, the membership assignment may have uncertainty attached and become instead a "fuzzy" set. For example, an F-16 may belong to the set of fighter jets with membership 0.90, also belong to the set of bombing aircraft with membership 0.6 and to the set of reconnaissance vehicles with membership 0.5. The uncertainty in defining set membership is represented by numerical values such as these. Unlike probabilities, there is no restriction for memberships to sum to one.

Fuzzy control system methods have been developed for analyzing complex systems and problems where the physical model or underlying processes are not known and where the input and output uncertainties may only be known in terms of language information (e.g., "nominal," "good," "bad," "poor," etc.) (Ross 1995). Relationships between outputs and inputs of a complex system are specified according to sets of rules and conditions. The rules, inputs and outputs accommodate uncertainties in information by using natural language terms, making it convenient for use by the experts providing information. In first strike stability, for example, the inputs and outputs would be the strikes and their results (targets killed), and the rules would be constructed describing the engagements. Fuzzy methods would provide the mechanisms to represent the many uncertainties involved in these inputs, outputs and rules.

Many decision problems contain much uncertainty in their structure and in the information available. They can be addressed using fuzzy cognitive maps (Kosko 1997). Fuzzy cognitive maps are a means of explaining political decision making processes, by combining the use of neural networks and fuzzy logic. Variable concepts are represented by nodes in a directed graph. The value at the node represents the degree in which that concept is active in the system at a given time. Qualitative simulation permits experimentation within the complex problem. Such analyses are done prior to gathering information (often expensive to obtain) for a more quantitative model. A detailed example modeling the onflict in Kosovo is found in Taylor (1999).

Wth the expanded use of fuzzy logic and set theory, other interpretations of uncertainty have emerged that differ from the foundations of probability theory. Zadeh (1996) proposed possibility theory. This theory provides methodology for addressing such questions as: how likely is it that China has strategic missiles of a certain type, or how large is the Chinese military force? Its axioms and properties are similar to probability, but they operate on the “possibility” and "plausibility" of events often in language (non-numeric) terms and are based on the concepts in fuzzy set theory.

Economists, engineers and computer scientists have applied these theories to problems in soft computing (Dubrois and Prade 1988), human cognition and knowledge management. While these theories have their axioms and principles of operation, such as coherence, they are broad enough to handle the large uncertainties inherent in human knowledge and information. And they are designed to more closely represent the cognitive processes which, too often, appear as irrational behavior and inconsistencies when attempting to apply game and probability theory.

3.7 What is Stability?

Several game theoretic notions of stability have been mentioned above in the context of the standard two-person game, ranging from actions that guarantee (statistically) a given payoff, saddlepoint solutions that define Nash equilibria for all individual players involved, and negotiated agreements that optimize some collective measure of joint benefit. Such notions have the benefit that stability is well defined in each case. As noted, however, none of these notions seems especially applicable to crisis stability or geopolitical stability.

Loosely speaking, stability measures involve many considerations (e.g., economic, political and military), many players (a situation that's stable for one player may not be stable for another), and should be capable of adapting to changes over time (a situation that's stable today may not be stable tomorrow). All of which raises the question: what is a stability measure? In other words, if a function *s(x)* of specified variables *x* were given, would we know if that function were a "stability measure" or not? And if so, how would we know it? The situation is very different from bench science, where it's often possible to conduct physical experiments which can, to some degree, confirm a hypothesis or validate a model.

From a historical standpoint, people have mused over stability for years. The subject has been of great interest since the 1960s, with much written about it, but there haven't been any widely accepted resolutions to the problem. Although "The nuclear world is full of multiple equilibria; there is more than one solution to the problem of deterrence and stability" (Weber 1991, p. 304), no one appears to have yet found a stability measure to fully capture any of them.

4. Force Stability: An Example

4.1 Background

We now consider a specific problem where a game theoretic formulation is helpful in assessing stability. That problem involves simulated nuclear exchanges. Force stability applies to a situation when two or more countries have reached an extreme crisis in which each believes the other may launch an attack. The situation is stable if the incentive for a country to strike first, rather than to plan to retaliate to the adversary's attack, is relatively weak. For this reason, force stability is also known as first strike stability. Unlike many stability settings, such as geopolitical stability, the action space for force stability problems is well defined and there is a solid basis for the expected payoffs and stochastic modeling.

In the most stable situation, it is possible for each side to inflict unacceptable damage upon retaliating against a first strike, so that neither side is motivated to initiate an exchange, thereby allowing the crisis to be resolved by other means. While force stability can be misleading when interpreted in a vacuum (between friendly countries, for example, who wouldn't attack each other even if their forces were of disparate military strength), it may have considerable relevance to future arms reduction agreements.

The computer simulation for nuclear exchanges postulates a war having multiple engagements, the term “multiple engagements” meaning that several attacks are to take place sequentially. Each side knows of the other’s weapons arsenals and so-called value assets (targets such as cities and industrial centers, which pose no immediate military threat but which are important to destroy). A scenario is specified, where one country strikes another, the other then retaliates (concluding the first engagement), and this is followed by additional engagements between the same or other countries. Obviously, each side wants to optimize its targeting, so as to do the most damage to its enemies while incurring minimal damage to itself. Combinatorically, the number of weapon-target strategies imbedded within the multiple engagements is finite, but too large for exhaustive calculation. Thus, optimization algorithms are necessary.

This modeling of warfare is not 100% accurate, in that intelligence regarding the other side’s arsenals and weapons performance is not perfect. Moreover, in real wars, each side does not wait patiently for the other side to finish its attack before responding. Nonetheless, much can be learned by examining strike and counterstrike strategies over a wide range of cases (e.g., a case with the U.S. and Russia and at their current arsenals, another with their projected arsenals under START agreements, and others related to so-called breakout scenarios where one side clandestinely develops weapons beyond those allowed by negotiated agreements).

The military simulation group at Los Alamos has developed the MESA/SM code (Multiple Engagements involving Strategic Arsenals, with Stability Metrics) for simulating multi-polar military exchanges (Anson and Stein 1999). It has the capability to optimize over repeated engagements involving multiple parties, albeit under deterministic conditions. Because of the detailed resolution of the code, producing strategies at the weapon/target level, considerable realism is achieved.

4.2 Stability Measures

Probably the most commonly used measures of force stability are Kent-Thaler indices (Kent and Thaler 1989). Kent-Thaler analysis, in the bi-polar scenario with a single engagement, calculates “costs” to each adversary in the event that they strike first and in the event that they are struck first and launch a retaliatory strike (each side is assumed to launch only once). Each side’s targets are divided into two categories: weapon targets and value assets, and it is assumed that the retaliatory strike attacks only value assets as there will be no further engagements. It is also assumed that each side makes use of all its weapons, so that the only optimization necessary is the first striker's decision of what fraction of weapons to aim at value assets.

Often, the two sides are assumed to have only one type of warhead, although this assumption is not critical. The amount of damage done by a given number of weapons is assumed to be deterministic, although typically attackers experience diminishing returns as they launch more weapons. A traditional cost function is then

$$C_S = f_S + \lambda(1 - f_O)$$

where C_S denotes the cost to one's self, f_S denotes the fraction of one's own value assets destroyed, f_O denotes the fraction of the opponent's value assets destroyed, and λ is a constant (often taken to be roughly 0.3). The minimum of C_S occurs, obviously, at zero, where a country loses none of its value assets while destroying all the assets of its opponent(s). The optimal first strike strategy is a combination of attacking an opponent's weapons (i.e., one warhead could destroy an opponent's platform, thereby preventing all of that platform's warheads from destroying one's own value assets) as well as attacking the opponent's value assets directly.

The cost C_S is computed for the two choices of "self" and the two possible first strikers, and then a stability index is

$$\frac{C_S(1,1)}{C_S(1,2)} \times \frac{C_S(2,2)}{C_S(2,1)}$$

where $C_S(i,j)$ denotes the cost to country *i* when country *j* strikes first. Low values of this index indicate that at least one of the countries may believe that they cannot afford to risk a first strike by the other.

Nyland (1998) modifies this analysis slightly, by assuming that the first striker attacks weapons up to the "point of diminishing returns" in the curve of opponent's warheads surviving as a function of one's own warheads launched at weapon targets. Nyland calculates stability indices using the number of warheads the second striker has available to launch. Another concept of stability, closest to the Mutually Assured Destruction doctrine, considers the situation stable if each side is assured of inflicting unacceptable damage in retaliation against a first strike.

As is apparent, such stability metrics are poorly suited to the stochastic nature of the problem. If, for example, a warhead-target combination is associated with a kill probability p_k, the damage inflicted by the warhead is stochastic. That is, the cost to each side of an exchange isn't known (exactly) in advance. Substituting statistical averages for the $C_S(i,j)$ can lead to misleading notions of stability, as we show later.

4.3 Simulated Nuclear Exchanges

4.3.1 Input Data

Extensive input files are needed to characterize the exchange. Each side's arsenal must be detailed, in terms of the number of platforms, missiles per platform, and warheads per missile. Value assets must also be described. Weapons platforms as well as value assets are legitimate targets, the former in order to preempt a counterstrike and the latter to cripple the enemy's infrastructure.

For each weapon/target combination, there is a "damage expectancy" (*DE*). The *DE* is a number between 0 and 1 which corresponds to the kill probability that a weapon's warhead will destroy its target. In exchanges with highly lethal weapons, the notion of an all-or-nothing target "kill" represents a reasonable modeling, in that partial damage is not a likely outcome. For warheads that behave independently, if *w* warheads are fired at a single target, the chance that all of them will miss is $(1\text{-}DE)^w$, and thus the probability that the target will be destroyed in the attack is $1 - (1\text{-}DE)^w$. Upon combining all such results, the possible outcomes of a first strike are numerous; and, because each subsequent engagement depends on the outcome of its predecessor, a wide variety of outcomes for a multiple engagement scenario is possible.

For optimization, a linearized version of the above expectation is useful and is incorporated into the current version of the MESA/SM code. We will refer to this assumption as the deterministic warhead assumption. That is, suppose that each warhead fired deterministically kills a fraction *DE* of its target. This interpretation can be computationally convenient, making determination of launch strategies possible by solving a linear programming problem with constraints rather than by using more time consuming integer routines.

The deterministic warhead assumption is reasonably accurate for high *DE* values, say, $DE \geq 0.9$, when there is an added constraint that damage may not exceed that which is physically possible (e.g., firing two warheads with DE = 0.9 at a single target does only one unit worth of damage, not 1.8 units worth). For smaller *DE* values, a nonnegligible bias is introduced by this assumption.

Suppose, for example, that an attacker has ten warheads, each with DE = 0.5, to launch at 5 targets. The deterministic warhead assumption leads to firing w = 2 warheads at each target and assuming that all targets will be killed. Stochastically, however, the chance of killing all 5 targets is $[1 - (1\text{-}DE)^2]^5$, here equal to 24%. It is clear that the deterministically expected damage is biased relative to the actual expected damage in such a way that the attacker believes that the attack will be more successful than usually is the case.

Still other issues arise when the stochastic nature of the exchange is considered. An example to illustrate the point: suppose there exists a single platform having 10 warheads, and it's fired on with DE = 0.8. Then, 80% of the time, the platform and its 10 warheads are destroyed, while 20% of the time, all 10 warheads remain. Basing a counterstrike on expected values assumes that there are 2 warheads remaining for retaliation, representing a statistical average with respect to the damage expectancy DE = 0.8. Of course, there are NEVER two warheads remaining, and any planned retaliation using 2 warheads may look very different than either actual counterstrike, with 0 or 10 warheads. Moreover, it is clear that substantial variability is introduced by the all-or-nothing attack on MIRVed targets, in that killing the targets offers substantial benefit to the attacking side, while not killing them creates the potential for much damage to be done upon retaliation.

In some domains, with single-warhead platforms, large numbers of weapons on both sides, and uniformly high DE values, deterministic expected value calculations are often reasonably close to their stochastic counterparts. With fewer weapons and greater MIRVing, as in the above example, the disparity between expected value assumptions and stochastic ones become greater. When the cascading effect is added, there can be considerable difference between stochastic and expected value results.

Cascading is the term that describes the collective effects of each engagement on the subsequent engagement(s). If, in the first engagement, one side is "lucky" (i.e., it emerges with less damage than was expected) and the other is unlucky, the former is positioned for greater success in subsequent engagements, while the latter is likely to incur greater damage than otherwise. The bottom-line effect of cascading can be large variability in costs for multi-engagement scenarios.

For all the above reasons (optimistic bias in deterministic assumptions, the role of damage assessment in planning, MIRVing of targets, and cascading), stochastic modeling has major implications for determining optimal strategies, and thus for first strike stability.

4.3.2 Optimal Stochastic Strategy

The lack of perfect predictability has an effect on defining optimality. In what follows, optimal strategies are defined in terms of expected values for each strike; that is, the best strategy, by definition, gives the best expected value response for the exchange. This is consistent with most current treatments of force stability. For completeness, however, it is noted that other strategies exist. As one example, if one side has a military advantage and is risk averse, it may want to minimize the probability of a negative outcome, even at the cost of reducing its expected margin of victory (similar to a sports team adopting a very conservative strategy when it has the lead in a game). Conversely, an underdog may choose to maximize its probability of a positive outcome even if the expected margin of defeat becomes larger (again, there are obvious examples from the world of sports). Such stochastic-based notions of stability have not been seriously examined in the force stability literature.

Conflicts begin with a set of goals for each country. These goals are of the form: in the event of a nuclear exchange with a certain country, it is desired to destroy some portion of that country's weapons and some portion of its value assets. Goals may be stratified, in the sense that some types of weapons/assets may have different priorities associated with their destruction. In principle these goals can model such concepts as countries' pain thresholds: they will target higher portions of their weapons at enemies' weapons if they are particularly resistant to absorbing damage to value targets.

Determination of each side's optimal strategy is, in real problems involving thousands of warheads and targets on each side, a huge computational task. For a single warhead with a high *DE* against a weapon target of the other side, a choice must be made between:

a) firing only a single warhead at a target (not "wasting" other warheads on a target that the first warhead is likely to kill) and firing another warhead at that target only if a subsequent damage assessment shows the target undestroyed, or

b) firing multiple warheads at a target to increase the probability of destruction, lest the target go undestroyed in the strike and then be fired on in an immediate retaliatory strike.

Hedging — deciding how may weapons to use in a first strike and how many to reserve for later strikes — is also important.

In determining the best strategies of all parties in a multiple engagement, the MESA/SM code carries out a complicated optimization based on genetic algorithms, in a process similar to dynamic programming. Details are beyond the scope of this report but can be found elsewhere (Anson and Stein 1999; Anandalingam and Friesz 1992). What can be said, though, is that the deterministic warhead assumption is used: the optimal retaliation is computed using the deterministically calculated number of weapons destroyed in the first strike, strategy for the second engagement is determined assuming that outcome for the first engagement, and so on.

To guide the optimization, quantitative metrics are established. Each side attempts to optimize its own situation, which involves incurring minimal damage to its own value assets while achieving its military objectives in destroying the opponents' value assets. The quantity C_S above is one such metric.

In the current MESA/SM implementation, a lengthy iterative algorithm is carried out to minimize the ultimate cost, from each side's standpoint, of its targeting strategy. Upon computing costs for each side under each scenario of interest (i.e., with each side assumed to strike first), stability can be assessed.

The role of perceptions in optimizing strategy is abundantly clear. Each side plans its own attack based on its perceptions. The two sets of perceptions may not be the same (e.g., one side's intelligence estimates of the other side's damage expectancies in striking certain targets may not match the values assumed by the other side). Indeed, the *DEs* for one's own attack are estimates based on somewhat limited data and may not be accurate to the 99^{th} decimal place. Depending on the magnitude of subjective uncertainty in evaluating *DEs*, there could be serious implications for the strategy selected.

4.4 A Simple Tri-polar Example

To illustrate the effect of stochastic behavior on stability metrics, we consider a simple tri-polar example. Input parameters are purely notional, in part to avoid potential classification issues with using values that reflect actual data. Moreover, the scope of the example is deliberately kept small (each side having at most two weapon types) so as not to introduce spurious complications.

Details are as follows. Three lesser powers are involved in a tri-polar nuclear exchange, a scenario that is not beyond the realm of possibility (for example: India, Pakistan, and China). In the hypothetical example considered here, Side #1 has MIRVed and non-MIRVed missiles, 10 of the former (with 10 warheads per missile) and 10 of the latter. Side #2 has 10 of the former and 40 of the latter, and Side #3 has only 20 non-MIRVED missiles. The numbers of value assets for the three sides is assumed effectively infinite relative to the limited arsenals.

The kill probabilities are modest because of an assumption that fledgling nuclear powers are not as proficient as the superpowers. Those kill probabilities, for each warhead and each of its potential targets, are as follows:

Targets

	Side #1 MIRVs	Side #1 non-MIRVs	Side #1 Assets		Side #2 MIRVs	Side #2 non-MIRVs	Side #2 Assets		Side #3 non-MIRVs	Side #3 Assets
Side #1 MIRVs					.25	.25	.25		.25	.25
Side #1 non-MIRVs					.40	.50	.80		.50	.80
Side #2 MIRVs	.25	.25	.25						.25	.25
Side #2 non-MIRVs	.40	.50	.80						.50	.80
Side #3 non-MIRVs	.40	.50	.80		.40	.50	.80			

In the simulated battle, the first engagement is between Sides #1 and #2 (Side #1 launching first and Side #2 then retaliating), the second engagement between Sides #3 and #1 (Side #3 launching first and Side #1 then retaliating), and the third engagement is between Sides #2 and #3 (Side #2 launching first and Side #3 then retaliating). The optimal strategy for each side involves so-called hedge quantities. That is, Side #1 will expend some of its weapons in the first strike, will then have some weapons destroyed in Side #2's counterattack, and will finally expend its remaining weapons in the next engagement. Determining an effective overall strategy involves using the right portion of its weapons in the first engagement and “hedging” the rest. Sides #2 and #3 must hedge in their first battles as well.

The MESA/SM code is run to obtain optimal targeting for this example. In a deterministic framework, each side believes it knows, based on the first strike, exactly how many targets remain. Through a complicated search process, the deterministic warhead assumption is propagated through each strike of each engagement. Converting the results to percentages for use in a stochastic simulation, the exchange proceeds as in Table 1.

Note that there is some artificial behavior in the above targeting scheme. In the first engagement, for example, Side #2 uses the knowledge that Side #1 will not attack it again. Thus, there is no reason for Side #2 to retaliate against Side #1’s weapons: those weapons would otherwise be used against Side #3, thereby making it less likely that Side #3’s weapons would later kill Side #2’s value assets. Moreover, in each side’s final attack of the war, it expends all its weapons because there is no point in saving them. Also, in the stochastic implementation, it is assumed that simply prorating the deterministic targeting strategy gives reasonable results when applied in stochastic mode where the number of weapons remaining after each attack is not (exactly) the same as anticipated deterministically; the more appealing approach of recomputing the targeting for each random outcome is not computationally feasible.

Upon propagating the deterministic warhead assumption through the successive engagements, the following value assets are killed:

Side #1: 2.8 Side #2: 2.6 Side #3: 21.9

These values reflect the initial conditions (Sides #1 and #2 started with more weapons than Side #3, and thus could destroy more of the other sides’ weapons before those weapons could be used on their own value assets).

It is straightforward to simulate the above war, using random numbers to decide whether certain targets are killed in the first strike, then simulating the retaliation to the first strike, and so on. Results from the stochastic simulation can then be compared to the deterministic results above. Results are shown in Figures 1 and 2.

Table 1

Attack Strategy for Hypothetical Example

First Engagement:

- First Attack: Side #1 → Side #2
  - 26% of MIRVs attack Side #2's MIRVs
  - 0% of MIRVs attack Side #2's non-MIRVs
  - 10% of MIRVs attack Side #2's value assets
  - 92% of non-MIRVs attack Side #2's MIRVs
  - 6% of non-MIRVs attack Side #2's non-MIRVs
  - 2% of non-MIRVs attack Side #2's value assets
- Second Attack: Side #2 → Side #1
  - 8% of non-MIRVs attack Side #1's value assets
  - 100% of MIRVs remaining attack Side #1's value assets

Second Engagement:

- First Attack: Side #3 → Side #1
  - 66% of non-MIRVs attack Side #1's MIRVs, if any remain
  - 1% of non-MIRVs attack Side #1's value assets
- Second Attack: Side #1 → Side #3
  - 100% of MIRVs remaining attack Side #3's value assets

Third Engagement:

- First Attack: Side #2 → Side #3
  - 35% of non-MIRVs remaining attack Side #3's non-MIRVs
  - 65% of non-MIRVs remaining attack Side #3's value assets
- Second Attack: Side #3 → Side #2
  - 100% of non-MIRVs remaining attack Side #2's value assets

Apparent in Figure 1 is the stochastic variability. There, a histogram of the value assets lost by Side #1 in the first strike is displayed. Over the 1000 simulated wars, assets lost range from 1 to 28, in contrast to the deterministic value 2.8 cited above. In other words, the nonlinearities are such that the expected (stochastic) result is very different than anticipated, and the variability around that average is nonnegligible. This large variability results from Side #1 allocating just enough warheads to destroy, under the deterministic warhead assumption, all of Side #2's MIRVs. Stochastically, however, one or more of Side #2's MIRVs survive with some probability, doing substantial damage in those cases. Indeed, such a stochastic distribution raises the notion of stability metrics based on probabilistic outcomes — e.g., a situation is stable if the probability of one side dominating the exchange is sufficiently small, or if the range of likely outcomes is sufficiently narrow.

A second simulation was run, reversing the roles of Sides #1 and #2 in the first engagement and using MESA/SM to derive the optimal targeting. In so doing, force stability can be examined. Similar to the discussion in Section 4.2, consider the stability metric

$$S = \frac{A(1,1)}{A(1,2)} \times \frac{A(2,2)}{A(2,1)},$$

where $A(i,j)$ denotes the assets lost by Side i (i=1, 2) when Side j strikes first in the first engagement. Based on the deterministic warhead assumption, propagating through the engagements gives the result S = (2.8 / 3.5) $\times$ (10.2 / 2.6) = 3.1.

Stochastically, a histogram of the simulated values of the stability metric is displayed in Figure 2. Note the wide range of values that is achievable (there is a lump of probability at $S = \infty$, in that a small fraction of simulated battles yield some $A(i,j)$ = 0). This wide range is partly due to using ratios to assess stability, which tends to exaggerate the stochastic effects in the assets destroyed. In any event, citing the average value of a stability metric with no mention of its associated variability can give a misleading picture.

Several variations on the above example could be considered. If Side #1 were risk averse, it could limit the range of damage to its value assets by firing a greater portion of its weapons at its opponents' weapons than in the above example. Conversely, it could adopt an aggressive strategy, launching only one warhead at each of its opponent's missiles and hoping to kill all of them in a very lucky (and unlikely) series of events; this would then allow for increased targeting of value assets. Such a high risk strategy would lead to a wider range of outcomes, depending on the degree of success in the first strike.

.

.

Other aspects of nuclear exchanges could also be drawn into this discussion. Damage assessment is important in a stochastic environment (why fire a second missile when the first has already killed the target?) but not in a deterministic world (since it is known, a priori, what damage will be done). Strategies involving damage assessment could be quite elaborate relative to the one considered above. Also, imperfect knowledge of damage expectancies, such as when each side must guess at the effectiveness of the other sides' weapons, introduces another source of uncertainty.

5. Conclusion

Stochastic elements play an important role in stability assessment. In force stability, for example, uncertainty enters in two ways: imperfect knowledge concerning an adversary's performance parameters (such as their damage expectancies) and the random outcomes of attacks (determined by kill probabilities p_k). These factors lead to a situation where the game theoretic payoff is unknown in advance.

There are two main implications of the stochastic behavior. The first is that an expected value approach produces misleading conclusions. Nonlinearities in the cascading effect are such that propagating expected value results through the multiple engagements does not usually lead to valid conclusions. Stochastic simulations are required to gain a good understanding of the dynamics.

The second problem with deterministic approaches is that the range of probable outcomes is not quantified. As seen in the example, random variability can be large. Strategies can be devised to minimize risk (these would be attractive to a side which has a decided advantage in the number of weapons) or to maximize risk (which might be attractive to a side which must take chances in order to prevail). As a consequence, the definition of "optimal" strategies is affected, with obvious implications for stability.

6. Acknowledgments

The authors thank Doug Anson and Myron Stein of the Los Alamos Military Simulation group for allowing us access to their MESA/SM code. We also thank James Scouras, whose presentation on force stability inspired some of the discussion given in Section 4.

Bibliography

Allison, G. T. (1971), *Essence of Decision: Explaining the Cuban Missile Crisis*, Boston: Little-Brown.

Anandalingam, G. and Friesz, T. L. (1992), "Heirarchical Optimization: An Introduction," *Annals of Operations Research* **34** 1–11.

Anson, D. and Stein, M. (1999), "MESA/SM: A New Multi-Polar Strategic Force Planning & Stability Methodology," Los Alamos National Laboratory technical report.

Booker, J.M and Bryson, M.C. (1985) "Decision Analysis in Project Management: An Overview," IEEE Transactions on Engineering Management, Vol. EM-32, p. 3-9.

Brams, S. J. (1985), *Superpower Games: Applying Game Theory to Superpower Conflict*, New Haven: Yale University Press.

Dresher, M. (1961), *Games of Strategy*, Englewood Cliffs: Prentice-Hall.

Dubrois, D. and Prade, H. (1988), *Possibility Theory: An Approach to Computerized Processing of Uncertainty*, New York: Plenum Press

Epstein, R. A. (1977), *The Theory of Gambling and Statistical Logic*, New York: Academic Press.

Fischoff, B., Slovic, P., and Lichtenstein, S. (1982), "Lay Foibles and Expert Fables in Judgments About Risks, *The American Statistician* **36** 240–255.

Giles, G. F., Doyle, J. E., Dunn, L. A., Gentry, L. W., and Sandrock, J. H. (1994), "Strategic Personality," SAIC Technical Report.

Howard, N. (1971), *Paradoxes of Rationality: Theory of Metagames and Political Behavior*, Cambridge: MIT Press.

Isaacs, R. (1965), *Differential Games*, Mineola, N.Y.: Dover.

Judd, B. R., Younker, L. W., Hannon, W. J., Strait, R. S., Sicherman, A., and Meagher, P. C. (1988), "Decision Framework for Evaluating Compliance with the Threshold Test Ban Treaty, Livermore National Laboratory report UCRL-53830.

Keeney, R. L. and Raiffa, H. (1993), *Decisions with Multiple Objectives: Preferences and Value Tradeoffs*, Cambridge: Cambridge University Press.

Kent, G. and Thaler, D. (1989), "First Strike Stability: A Methodology for Evaluating Strategic Forces," RAND Corporation Report R-3765-AF.

Knowledge Transfer International (1999) web site for terminology: http://ktic.com/topic6/13_TERM2.HTM, and University of Calgary web site: http://sern.ucalgary.ca/KSI/KAW/KAW99/KAW99Call.html.

Kosko, B. (1997) *Fuzzy Engineering*, Upper Saddle River, N.J.: Prentice-Hall.

Krakowski, R. A. (1999), "A Multi-Attribute Utility Approach to Generating Proliferation Risk Metrics," Los Alamos National Laboratory technical report LA-UR-96-3620.

Meyer, M. and Booker, J. (1991), *Eliciting and Analyzing Expert Judgment: A Practical Guide*, London: Academic Press.

Nyland, F. (1998), "The Index of Warhead First Strike Stability," R-128-ACDA, U. S. Arms Control and Disarmament Agency.

Ross, T.J. (1995), *Fuzzy Logic with Engineering Applications*, New York: McGraw-Hill

Schelling, T. C. (1960), *The Strategy of Conflict*, Cambridge: Harvard University Press.

Shubik, M. (1975), *Games for Society, Business, and War*, New York: Elsivier.

Slovic, P., Fischhoff, B. and Lichtenstien, S (1982), "Facts versus Fears: Understanding Perceived Risk," p. 463–492 in *Judgment Under Uncertainty: Heuristics and Biases*, (Kahneman, Slovic and Tversky, editors), Cambridge: Cambridge University Press.

Straffin, P. D. (1993), *Game Theory and Strategy*, Washington, D.C.: Mathematical Association of America

Taylor, W. (1999), *Predicting...the Future in Kosovo, the Past in Vietnam*, to be published and web site: http://users.neca.com/williamtaylor/kosovo.html.

Von Neumann, J. and Morgenstern, O. (1944), *Theory of Games and Economic Behavior*, Princeton: Princeton University Press.

Weber, S. (1991), *Cooperation and Discord in U.S.-Soviet Arms Control*, Princeton: Princeton University Press.

Weirich, P. (1998), *Equilibrium and Rationality: Game Theory Revised by Decision Rules*, New York: Cambridge University Press.

Whitman, M. (1985) "Black Hats, Economists and Societal Risk Assessment," p. 1–8 in *Risk Assessment and Management*, (Lave, editor), New York: Plenum Press.

Zadeh, L. A. (1965) *Fuzzy Sets, Information and Control*, Vol. 8, p. 338-353.

Zadeh, L. A. (1996), Possibility Theory and Soft Data Analysis, in Klir, G.J. and Yuan, B. (editors) *Fuzzy Sets*, *Fuzzy Logic*, *and Fuzzy Systems*, River Edge, N.J.: World Scientific Press, Vol 6, p. 481-541.

Zagare, F. C. and Kilgour, D. M. (1995), “Assessing Competing Defense Postures,” *World Politics* **47** 373–417.