Steven J. Brams / D. Marc Kilgour

Is Nuclear Deterrence Rational, and Will Star Wars Help?

Abstract: Deterrence means threatening to retaliate against an attack in order to deter it in the first place. The central problem with a policy of deterrence is that the threat of retaliation may not be credible if retaliation leads to a worse outcome - perhaps a nuclear holocaust - than a side would suffer from absorbing a limited first strike and not retaliating. - The optimality of deterrence is analyzed by means of a Deterrence Game based on Chicken, in which each player chooses a probability (or level) of preemption, and of retaliation if preempted. The Nash equilibria, or stable outcomes, in this game are compared with those in a Star Wars Game, in which the preemption and retaliation levels are constrained by the defensive capabilities of each side. Unlike threats in the Deterrence Game, which can always stabilize the cooperative outcome, mutual preemption emerges as an equilibrium in the Star Wars Game, underscoring the problem - particularly if defensive capabilities are unbalanced - that deterrence will be subverted by the development of Star Wars.

Nuclear deterrence is the cornerstone of the national-security policies of not only the superpowers but other nations as well. By threatening untoward action against an opponent who initiates conflict, even at great potential cost to oneself, one seeks to deter the opponent from committing aggression in the first place.

The controversy over the viability of nuclear deterrence has largely concerned the rationality of adhering to a policy that can lead to enormous destruction - perhaps even mutual annihilation - if the policy fails. The party attacked would seem foolhardy to bring upon itself a disastrous outcome if, by compromising or - heaven forbid! - capitulating, it could do better. On the other hand, by fighting (irrationally?) to the bitter end, it would seem to violate the very canons of rationality on which deterrence rests. Yet by caving in, or indicating that it might, it would seem to invite attack.

A number of different nuclear doctrines to support deterrence have been proposed, perhaps the most notable being MAD, or 'mutual assured

Analyse & Kritik 9 (1987), S. 62–74

destruction'. The inclusion of mutual in the MAD doctrine implies that each side can destroy the other, even if attacked first; this reciprocal vulnerability is presumed to make deterrence stable, at least as long as the mutual destruction is 'assured'.

Sometimes MAD is used to denote 'mutual assured deterrence', with the means for assuring deterrence not necessarily assumed to be the destruction of society. Other terminology is less lurid than MAD. 'Countervalue', which is stressed in the doctrine of mutual assured destruction, refers to the destruction of cities and industries, whereas 'counterforce' stresses the destruction of military forces, particularly missile sites, and command and control facilities. Still different strategies such as 'damage-limitation' and 'war-fighting' defenses after a limited nuclear attack - should deterrence fail - are also discussed in the national-security literature and are now part of the nuclear vernacular.

The rather arcane debate about nuclear deterrence and its alternatives is generally not about whether one should respond to attack, but how. Here our concern is broader - with the nature of deterrence itself: the conditions under which one should respond to an attack, with what degree of certainty, and at what level. Our purpose is not so much to describe optimal threats to deter an opponent - though such prescriptions will come out of the models we describe - but rather to show that deterrence is amenable to rational analysis. The foundation of this analysis is the mathematical theory of games, whose application to the problem of deterrence helps to clarify the main strategic issues.

Before attempting to apply this theory directly, consider what general argument can be used to justify deterrence, assuming that it is costly for a threatener to carry out a threat if attacked. While conceding that it is irrational to carry out a threat in a single play of a game, one might argue that it may well be rational in repeated play (Brams/Hessel 1984). The reason is that a carried-out threat enhances one's credibility - in doing the apparently irrational thing in a single play - so that, over the long run, one can develop a sufficiently fearsome reputation to deter future opponents. Thereby, although losing on occasion in the short run, one can gain over time.

British Prime Minister Margaret Thatcher evidently made this calculation when she responded to Argentina's invasion of the Falkland Islands in 1982 by dispatching the British fleet. The conflict was very damaging to both sides, but Britain's successful invasion left little doubt about that country's resolve in future territorial disputes, such as might occur over Gibraltar.

This is not a satisfactory argument, however, if carrying out a threat leads to something as unthinkable and irredeemable as nuclear war between the superpowers, which probably would occur only once. Unlike deterrence in a conventional conflict - between, say, a superpower and a smaller country without nuclear weapons, in which one's willingness to carry out a threat will affect one's reputation in future conflicts - credibility in a game without a sequel is purely academic.

To model deterrence between two nuclear powers, we begin with the two-person game of Chicken, in which each player can choose between two strategies: cooperate (C) and do not cooperate ($\bar{C}$), which in the context of deterrence may be thought of as 'do not attack' and 'attack', respectively. These strategies lead to four possible outcomes, which the players are assumed to rank from best (4) to worst (1). These rankings are shown as ordered pairs in the outcome matrix of Figure 1, with the first number indicating the rank assigned by the row player (called "Row"), and the second number indicating the rank assigned by the column player (called "Column"). Chicken is defined by the following outcome rankings of the two players:

1. Both players cooperate (CC) - next-best outcome for both players: (3,3).
2. One player cooperates and the other does not ($C\bar{C}$ and $\bar{C}C$) - best outcome for the player who does not cooperate and next-worst for the player who does: (2,4) and (4,2).
3. Both players do not cooperate ($\bar{C}\bar{C}$) - worst outcome for both players: (1,1).

Outcomes (2,4) and (4,2) in Figure 1 are circled to indicate that they are <u>Nash equilibria</u>: neither player (Row or Column) would have an incentive unilaterally to depart from these outcomes because he would do worse if he did. For example, from (2,4) Row would do worse if he moved to (1,1), and Column would do worse if he moved to (3,3). By contrast, from (3,3) Row would do better if he moved to (4,2), and Column would do better if he moved to (2,4).

FIGURE 1

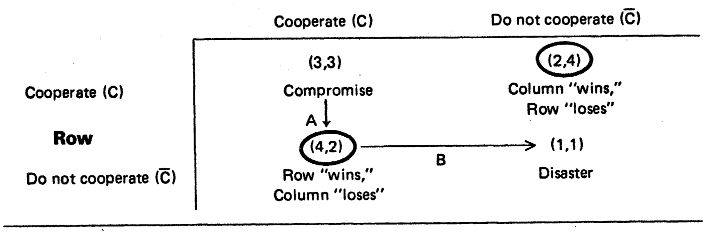

The shorthand verbal descriptions given for each outcome in Figure 1 suggest the vexing problem the players face in choosing between C and $\bar{C}$: by choosing $\bar{C}$, each can 'win' but risks disaster; by choosing C, each might benefit from compromise but could also 'lose'. Each Nash equilibrium shown in Figure 1 favors one player over the other, but the stability of these equilibria as such says nothing about which of the two - if either - will be chosen.

Although the (3,3) compromise outcome is the obvious candidate for the players to agree on, its instability would seem to rule it out as a durable solution. At some point each player might be tempted to depart from it to 'win', or at least threaten the other player with preemption.

One effect of threats in Chicken is not hard to grasp. If, say, Row threatens Column with the choice of $\bar{C}$, and this threat is regarded as credible, Column's best response is C, leading to (4,2), an apparent win for Row and loss for Column.

Clearly, the player with the credible threat - if there is one - can force the other player to back down in order to avoid (1,1). Although Row would 'win' in this case by getting his best outcome, Column would not 'lose' in the usual sense by getting his worst outcome but rather his next-worst.

This fact illustrates that Chicken is not a constant-sum game, in which what one player wins the other loses. That is why we have put 'win' and 'lose' in quotation marks here and in Figure 1. In variable-sum games like Chicken, the sum of the players' payoffs at each outcome (if measured cardinally by utilities rather than ordinally by ranks) is not constant but variable. This means that both players may do better at some outcomes (for example, (3,3)) than at others (for example, (1,1)).

The Deterrence Game is based on Chicken but adds two refinements: (i) the players can make quantitative choices of levels of cooperation (or non-cooperation), not just qualitative choices of C or $\bar{C}$; (ii) once these initial choices, which can be interpreted as levels of nonpreemption (or preemption), are made, the less preemptive player may choose a subsequent level or retaliation (Brams/Kilgour 1985a; 1985b).

To illustrate play in this game, consider the extreme case in which Row chooses maximum preemption and Column chooses no preemption initially. Thus, if the game starts out at (3,3), Row's preemption moves it to (4,2), as shown by arrow A in Figure 1. If we assume that the players have complete information about each other's initial choices, there would be no doubt that Row preempted; Column could then retaliate.

Now here comes the rub for Column, the player who cooperated initially. If he responds to Row's preemption and moves the game to (1,1), as illustrated by arrow B, he succeeds not only in punishing Row but also himself. This is precisely what makes retaliation in the Deterrence Game problematic: it would be better for Column to capitulate, accepting his next-worst outcome (4,2), than avenge Row's preemption by moving the game to (1,1), the disastrous outcome for both players.

True, if revenge in this situation is valued more highly than defeat by Column, there is nothing irrational about Column's choosing retribution against Row. But this choice is incompatible with the Figure 1 payoffs; moreover, revenge might be especially hard for a superpower to justify after suffering a limited nuclear attack that the vast majority of its population survives or, more realistically, a large-scale conventional attack, such as by the Soviet Union in Western Europe. Nuclear reprisal, after all, would almost surely result in a full-scale nuclear exchange, whose consequence might well be a nuclear winter in which everybody perishes.

If the rankings of outcomes in Figure 1 accurately describe the problem of retaliation in a nuclear confrontation between the superpowers, how can a policy of deterrence be justified that poses the threat of nuclear war and perhaps mutual annihilation? The superpowers in effect circumvent the problem of rationally responding to a first strike by irrevocably pre-committing themselves to retaliate if attacked (Brams 1985). Thereby they preclude themselves from making conciliatory choices at the very point at which it might be prudent to step back from the precipice.

In fact, command and control procedures that both superpowers now have in place specify preselected targets that will be hit once a first strike of a particular magnitude is detected. Even if the president is incapacitated, authority for the launching of a retaliatory strike devolves (to lower levels of command) to ensure that such a strike will actually be carried out (Bracken 1983; Blair 1985; Ford 1985; Lebow 1987; Carter/Steinbruner/Zraket 1987).

All this smacks of a 'doomsday machine', which responds independently of human decisionmakers. This is an exaggeration, of course, but it probably is accurate to speak of a 'probabilistic doomsday machine' (PDM) - one with built-in uncertainties due to possible failures in C^3I (command, communication, control, and intelligence), including the lack of will of political decisionmakers to order a second strike as well as a variety of technical problems that might arise.

Is a PDM sufficient to deter a first strike by an opponent? In principle, this will depend on whether the opponent thinks he can do better by attacking or by not attacking. Assume, for illustration, that the ranks in Figure 1 are cardinal utilities, or actual values, that each player associates with the four outcomes. If p is the probability that the PDM will function properly when one player - say, Row - attacks the other, then Column can deter Row if the payoff that Row obtains from not attacking, 3, is greater than the expected payoff he obtains from attacking. In this case, Row will obtain 4 with probability 1-p (PDM does not work so Column will not retaliate) and 1 with probability p (Column will retaliate). This calculation can be expressed by inequality $3 > 1p + 4(1-p)$, which is equivalent to $p > 1/3$.

In other words, Row will be well advised not to attack if Column's PDM has a greater than 1/3 chance of triggering retaliation. If the consequences of retaliation were much worse than 1 (for example, some large negative value, which might be the case in a nuclear conflict), only a very small probability p of retaliation would be required to make a first strike unprofitable if not perilous for the attacker.

Patently, certain retaliation is not necessary to deter an opponent in the Deterrence Game, at least one who makes the kind of expected-payoff calculation we have illustrated. In international conflicts, especially those that might involve nuclear weapons, there is abundant evidence that national decisionmakers are not reckless but, in fact, rather conservative in their choice of means to satisfy their goals.

Of course, we may not share their goals or even sympathize with them. This fact, however, is not consequential if we have a fairly good idea of what their goals are and, specifically, how they evaluate the possible outcomes that may occur. Given that they rank outcomes as in Figure 1, there will be some p less than 1 that will render the expected payoff obtained from attack (and subsequent retaliation) less than that obtained from not attacking.

This calculus, nevertheless, does not gainsay that retaliation is always costly to the player attacked in the Deterrence Game. This is why, to make his threat of retaliation credible, he must <u>precommit</u> himself to retaliate with a p above the threshold value we have illustrated.

The superpowers have made themselves credible by, in effect, constructing PDMs - somewhat beyond the control even of their top leaders - who might, conceivably, prefer to surrender in a crisis rather than retaliate against a first strike. The fact that they may not be able to countermand the PDM ensures that precommitments to retaliate are credible.

Probabilistic threats of retaliation that deter an attack will presumably depend on the level of the attack. As the level increases and the first strike brings the attacker closer to his highest payoff (before retaliation), the retaliator will have to increase his level of retaliation in order to reduce the attacker's payoff to an amount below what the attacker would get if he had not attacked in the first place.

Visually, one might think of the Deterrence Game as played on a square board, whose four corners give the payoffs shown for Chicken in Figure 1. If Row attacks by moving the outcome vertically from (3,3) toward - but not necessarily reaching - (4,2), Column can respond by moving horizontally from left to right, closer to (1,1). This will decrease Row's payoff, so that given Row's level of attack, at some point on the board defined by Column's level of response, Row will be indifferent between attacking and not attacking. Retaliation that carries the outcome to the right of this point, closer to (1,1), will definitely be worse for Row than staying at (3,3).

We assume in the Deterrence Game that the players' payoffs vary continuously as a function of the distances from the four corners of the board. Each can deter his opponent by threatening retaliation at some level greater than that which causes indifference.

We have calculated, in a variation of the Deterrence Game called the Threat Game, the minimal levels of retaliation that are required to deter attacks and have discovered that a policy of tit-for-tat reprisals may not be the best deterrent (Brams/Kilgour 1987). In many cases, a more-than-proportionate response is optimal against relatively minor aggression, a less-than-proportionate response against relatively major aggression. The precise levels - and the threshold at which 'more' becomes 'less' - depend on the payoffs of the underlying game of Chicken.

Historically, we will never know whether a strong policy of resistance against Hitler's early incursions would have prevented World War II. Nor can we predict that, after a limited nuclear first strike by one superpower, a diminished response on the part of the other will prevent World War III. There are, nonetheless, good rational reasons to believe that an effective deterrent may be one in which the level of retaliation is tailored more or less - but not strictly - to the level of aggression. By hitting relatively hard when the provocation is small, and backing off somewhat when large-scale conflict might prove catastrophic, one may at the same time discourage 'salami tactics' and defuse all-out escalation should deterrence fail.

These results support a modified tit-for-tat policy or, in the parlance of the U.S. Defense Department, "flexible response" or "graduated deterrence". But our results are more precise than these qualitative doctrines; they provide quantitative guidelines of the punishment that should be threatened in relation to the level of attack (Brams/Kilgour 1987). Specifically, as aggression increases, retaliation should also increase, but at a decreasing rate. Whether the threshold retaliation need be more than the aggression at low levels depends on the specific payoffs in the game, but at high levels one need never threaten retaliation commensurate with the provocation to deter it. The reason is simple: such a policy will move the game toward the mutually worst (1,1) outcome, which may be far worse than the (3,3) outcome; rational deterrence can always be achieved without the threat of such 'overkill'.

The (3,3) outcome is in fact a Nash equilibrium in the Deterrence Game and Threat Game when backed up by threats above the minimum level necessary to deter. In other words, the cooperative outcome in Chicken, which is not in equilibrium, can be stabilized by threats - at least if they are considered credible by an opponent.

What helps make them credible is that the threatener does not suffer unnecessarily great damage in carrying them out, making his precommitment to such retaliation more plausible. Although the threatened retaliation probably should be somewhat above the minimum level to guard against possible misperceptions or miscalculations by an opponent, it should never lead to complete devastation of the threatener. Otherwise it would appear incredible; the potential aggressor, suspecting such retaliation would never be carried out, might attack on this very presumption.

It is unfortunate, perhaps, that the palpable fear of annihilation rather than simple good will has prevented nuclear war for forty years. Yet good will alone is insufficient to sustain (3,3) in Chicken precisely because it is rational to defect from this cooperative outcome. Threats in the Threat Game - sometimes entailing retaliation greater than small-scale aggression, but always diminishing in proportion to increasing aggression - can, however, render deterrence rational.

The superpowers have flirted with their own extinction to make deterrence work. Ronald Reagan's SDI (Strategic Defense Initiative, or 'Star Wars'), which purports to provide a defense against nuclear weapons and avert an apocalypse, may eventually replace PDMs, which he finds unsavory.

Unsavory as nuclear deterrence may seem, we have argued that it can be grounded in game-theoretic rationality, given threats of retaliation that have a sufficiently high probability of being carried out. We are not, by the way, suggesting that the throw of dice or the spin of a roulette wheel should determine whether an American president or a Soviet general secretary will retaliate against a first strike but rather that the uncertainties of retaliation are already inherent in C^3I - and its possible breakdown.

Thus, a PDM, in substantial part devoid of human intervention, would appear to be a rational mechanism for stabilizing deterrence. But when it is shorn of its human element and rests so heavily on impersonal detection devices, computers, and the like, there would appear to be something inhumane and even morally repugnant in threatening horrendous destruction in order to deter a first strike. Should millions of innocent civilians be held hostage to maintain the proverbial 'delicate balance of terror'?

Star Wars holds out the promise of forestalling a preemptive strike by preventing many first-strike weapons from getting through, using one or more shields. If this attack can be stopped or largely blunted, then presumably the potential preemptor will think twice about attacking in the first place. Moreover, even if he does attack, his attack will not be nearly so effective as it would in the absence of a missile defense, thereby decreasing the value of striking first.

This argument for a strategic defense does not hinge on its being totally impenetrable, or 'leakproof', but instead on its lowering the expected payoff to an attacker of a first strike. Consequently, deterrence will be enhanced, which is today the primary justification the Reagan administration uses for Star Wars in light of the apparent impossibility of building a leakproof defense, at least in the foreseeable future.

However, the other side of the coin is that, with a Star Wars defense, each side will be able to degrade the effectiveness of a (retaliatory) second strike. This degradation will be especially upsetting if each side can be crippled or seriously damaged by a first strike, diminishing greatly its capacity to retaliate and thereby undermining deterrence. If a Star Wars defense is in fact possible, the key question is: Will the enhancement of deterrence, by making a first strike more uncertain, be offset by the undermining of deterrence because one's capability to retaliate, particularly after a devastating first strike, will be undercut?

We ignore here the enormous costs of building a Star Wars system. Our focus is solely on the strategic effects of Star Wars on deterrence, assuming that deterrence in some form will not be abandoned, at least until Star Wars is perfected. Yet the perfection of Star Wars, to the point that it becomes a leakproof system, is surely an extremely remote possibility.

To analyze the deterrence-enhancing versus the deterrence-undermining effects of Star Wars, we assume that Star Wars puts limits on the maximum first and second strikes of each player in the Deterrence Game (Brams/Kilgour 1986). That is, we introduce as new parameters in this game constraints on how far, say, Row can shift the outcome from (3,3) to (4,2) in a maximal first strike, and, in turn, how far Column, after suffering a first strike, can shift the outcome from (4,2) - or wherever the game is after the first strike - toward (1,1) and full retaliation.

We posit three different scenarios that assume different functional relationships between each side's first and second strike defense. Then we let these defenses vary from no defense to perfect defense, subject to these relationships. We will not try to describe the different scenarios here but instead will summarize our principal results, based on all the scenarios.

Generally, we find that, for low levels of strategic defense, deterrence can be maintained. The reason is that each side's threat of retaliation is still sufficient to deter an opponent from preemption, but as defenses improve this threat loses its force and the stability of (3,3) in the Deterrence Game is jeopardized.

At a calculable threshold value of defense, deterrence breaks down and it becomes rational for each side to attack the other. Not only can neither side be deterred by the threat of retaliation when its defense is sufficiently strong, but it also does better attacking preemptively than retaliating after being attacked.

This is a disturbing development, for it renders mutual preemption a Nash equilibrium in the Deterrence Game with Star Wars or simply the Star Wars Game; this outcome is never stable in the absence of Star Wars. (True, unilateral preemption to (4,2) and (2,4) are also Nash equilibria in the Deterrence Game, but given precommitted threats by both sides above the threshold level calculated earlier, they are dominated by the choice of (3,3).) In the Star Wars Game, by contrast, both sides may find it advantageous to attack each other simultaneously because, if each has a strong enough defense, neither side's threats of retaliation will be sufficient to deter an opponent.

Actually, an equality or near equality in the defenses of the two sides retards mutual preemption, whereas an imbalance in defenses aggravates it. For if one side has a much stronger Star Wars defense than the other, by attacking first it might be able to so weaken its opponent that it can effectively stop whatever retaliation the opponent can throw back. But the opponent can make this calculation, too, and realize that it would do better attacking itself - given it is about to be preempted - resulting in mutual preemption. Such preemption may be arrested either by credible threats of retaliation or, if less than credible, a mutual realization by the players that not attacking is still better than attacking with strong but not necessarily impenetrable defenses.

In our different scenarios, both mutual preemption and deterrence, as well as unilateral preemption, are Nash equilibria for certain levels of defense;

conditions under which one equilibrium may dominate another when they coexist are investigated. Perhaps the greatest peril occurs when there is no deterrence equlibrium. Then an extreme form of crisis instability may grip the players and lead them to an abyss. More probable in superpower relations, though, is that deterrence will remain reasonably secure, mainly because both sides have largely invulnerable second-strike capabilities (principally, submarine-launched ballistic missiles and cruise missiles) that Star Wars will have no effect on, at least presently.

At some point, however, perhaps in a severe crisis, crisis stability could be upset and preemption, perhaps even mutual preemption, might appear attractive. This has occurred at lower levels of superpower conflict, usually through surrogates, in different parts of the world. If we are to steer clear of nuclear preemption as a rational option, it is imperative that the superpowers recognize that they must carefully chart a course of balanced development of Star Wars defenses - if these ever become feasible - to avoid creating major instabilities, particularly in the period of transition from deterrence to defense.

The replacement of a deterrent policy depending on PDMs by a defensive posture grounded in Star Wars is not imminent. Until it occurs, it behooves us to understand the logic of nuclear deterrence and to improve upon it through calculations that make it as robust as possible. Star Wars, as we have modeled it, seems mostly an assault on this strategic logic.

Note

* The first part of this paper on deterrence is a slightly revised version of Brams/Kilgour (1986b), but the second part on Star Wars is new. The material on deterrence is based in part on Brams (1985, ch.1) and the several papers on deterrence and threats cited in the text. The models in these papers and others on escalation in an arms race, crisis stability, verification of arms-control agreements, and winding down if deterrence fails are developed within a common framework in Brams/Kilgour (1988). Brams gratefully acknowledges the financial support of the National Science Foundation under Grant No. SES 84-08505. Kilgour gratefully acknowledges the financial support of the Natural Sciences and Engineering Research Council of Canada under grant No. A8974.

Bibliography

Blair, G. (1985), Strategic Command and Control: Redefining the Nuclear Threat, Washington/D.C.

Bracken, P. (1983), The Command and Control of Nuclear Forces, New Haven/C.T.

Brams, S.J. (1985), Superpower Games: Applying Game Theory to Superpower Conflict, New Haven/C.T.

Brams, S.J./M.P. Hessel (1984), Threat Power in Sequential Games, in: International Studies Quarterly 28, 15-36

Brams, S.J./D.M. Kilgour (1985a), Optimal Deterrence, in: Social Philosophy & Policy 3, 118-135 (Reprinted in: E.F. Paul et al. (eds.), Nuclear Rights/Nuclear Wrongs, London 1986, pp. 118-135; and N. Garver/H. Hare (eds.), Naturalism and Rationality, Buffalo/N.Y. 1986, 241-262)

- (1985b), The Path to Stable Deterrence, in: U. Luterbacher/M.D. Ward (eds.), Dynamic Models of International Conflict, Boulder/C.O., 11-25

- (1986a), Deterrence Versus Defense: A Game-Theoretic Model of Star Wars, Mimeographed

- (1986b), Is Nuclear Deterrence Rational?, in: PS 19, 645-651

- (1987), Optimal Threats, in: Operations Research 35, forthcoming

- (1988), Game Theory and National Security, forthcoming

Carter, A.B./J.D. Steinbruner/C.A. Zraket (eds.) (1987), Managing Nuclear Operations, Washington/D.C.

Ford, D. (1985), The Button: The Pentagon's Command and Control System, New York

Lebow, R.N. (1987), Nuclear Crisis Management: A Dangerous Illusion, Ithaca/N.Y.

<u>Further recent work by Steven J. Brams and D. Marc Kilgour not cited here includes:</u>

- Rational Deescalation, in: D. Farmer et al. (eds.), Evolution, Games, and Learning: Models for Adaptation in Machines and Nature, Physica 22D, Amsterdam 1986, 337-350

- Notes on Arms-Control Verification: A Game-Theoretic Analysis, in: R. Avenhaus et al. (eds.), Modelling and Analysis in Arms Control, Berlin 1986, 409-419

- Winding Down if Preemption or Escalation Occurs: A Game-Theoretic Analysis, in: Journal of Conflict Resolution 31, December 1987, forthcoming

- Threat Escalation and Crisis Stability: A Game-Theoretic Analysis, in: American Political Science Review 81, September 1987, forthcoming

- Verification and Stability: A Game-Theoretic Analysis, in: A. Din (ed.), Arms and Artificial Intelligence, Stockholm, International Peace Research Institute 1987, forthcoming

- R. Avenhaus/J. Fichtner, The Probability of Nuclear War, Mimeographed 1986