A Continuous-Outcome Expected Utility Theory of War

JAMES D. MORROW
Department of Political Science
University of Michigan

A continuous-outcome expected utility model is presented that generalizes the expected utility theory of Bueno de Mesquita. An examination of the more general model uncovers several unstated assumptions within and produces new conclusions from, while supporting the basic logic of, the expected utility theory. Among the new conclusions is the finding that nations shifting their level of acceptable outcomes to a conflict upward or downward after fighting starts is perfectly consistent with a rational model. The derivations demonstrate the value of theoretical articulation, a task too often neglected in quantitative international relations, and provide a sound logical basis for the construction of systemic theories based upon the expected utility theory.

> "You know you never defeated us on the battlefield," said the American colonel.
>
> The North Vietnamese colonel pondered this remark a moment. "That may be so," he replied, "but it is also irrelevant."
>
> Conversation in Hanoi, April 1975 [Summers, 1981]

War in many ways is a game of chance. At the start of hostilities, neither side knows what the final outcome will be or how long it will take to reach that outcome. It is this uncertainty about the final outcome that allows each side to believe that it will benefit more by fighting than negotiating (Blainey, 1973; Wittman, 1979). Furthermore, wars are rarely resolved by the complete military destruction of one side; rather the fighting progresses to a point where both sides agree on a negotiated solution because they now believe it will be impossible to achieve all of their

AUTHOR'S NOTE: I would like to thank the referees for their comments on this article. I would also like to thank the Political Science Department at Michigan State University for intellectual and material support during the writing of this article.

JOURNAL OF CONFLICT RESOLUTION, Vol. 29 No. 3, September 1985 473-502

aims. The defender demonstrates sufficient military competence to dissuade the attacker from continuing the war. To be accurate in modeling decisions to go to war, we must model the diversity of outcomes and the actors' valuations of those outcomes to capture the essence of the decision each side faces.

Existing rational choice models of war decisions do not capture this diversity of outcomes. They view crisis and war decisions either as expected utility calculations over the discrete outcomes of winning and losing (Wittman, 1979; Bueno de Mesquita, 1981a) or as two-by-two games over the choice of fighting or not fighting (Snyder and Diesing, 1977). These models are inherently unable capture the variety of outcomes that war can produce, nor the effects of that variety on the decision to go to war. This article will present an expected utility model of war decisions based on a continuous spectrum of outcomes, show under what conditions that model can be reduced to a simple discrete expected utility model and the implications of that reduction, and derive several interesting new propositions about war and its occurrence. This model can be seen as a generalization of Bueno de Mesquita's expected utility model that eliminates certain logical problems with his approach (Zagare, 1982; Wagner, 1984; Majeski and Sylvan, 1984) and preserves the logically sound deductions while promoting a more general expected utility model. For game theoretic approaches, this model will allow a proper specification of the payoffs in the game for conflict outcomes. By correctly specifying the payoffs, greater accuracy in these models can be achieved.

This article then is an exercise in theoretical articulation, a task that is too rarely attended to in the field of quantitative international relations. Theoretical articulation is the clarification and extension of the theoretical structure of an existing paradigm. Clarification of the theoretical structure enhances the ease of application and instruction of a paradigm and creates a stronger structure to expand; extension of the theoretical structure demonstrates the ability (or inability) of a paradigm to predict, and thus explain, previously inexplicable phenomena. In addition, theoretical development of a paradigm allows the scientific community to test the logical validity and generality of the theories upon which they base their explanations. The successful articulation of a theory can accomplish four different tasks: one, flush out unstated assumptions; two, streamline the logic of the theory; three, direct additional tests of the theory; and four, expand competing theories to allow for their reconciliation. The last task is particularly important for fields where there is not an accepted theory, such as quantitative international relations. In a field where there is little agreement on the accumulated body of knowledge, explanations of existing results, rather than additional results, are needed. Such explanations require both the statement and articulation of theories. In the interest of furthering the accumulation of knowledge, I turn to the articulation of a particular theory, Bueno de Mesquita's (1981a) expected utility theory of war.

This article will first present the assumptions and logical structure of the expected utility calculation examined herein. Conclusions about the initiation of conflict will be drawn from the model. Conditions under which the entire calculation can be reduced to a simple discrete form will be deduced and the implications of those conditions will be discussed. One of these conditions will be a multilateral decomposition of the calculation that will allow the impact of any one third party on the conflict to be analyzed. The results will be illustrated with some simple examples, and several suggestive propositions about conflict will be drawn out. The logical soundness of both Bueno de Mesquita's discrete calculation and published criticisms of it will be examined, and both will be found wanting. The implications of that theory for empirical tests of expected utility propositions will be considered. Finally, the implications of the generality of the theory presented here for other theories of war will be presented in order to stress the importance of developing theories that are built on a sound logical basis.

ASSUMPTIONS OF THE THEORY

According to expected utility theory, an actor's expected utility provides a measure of his preferences for different courses of action (Luce and Raiffa, 1957; Riker and Ordeshook, 1973). If an actor prefers a particular (risky) course of action, then the expected utility for that course of action must be larger than any other course of action available. To calculate an actor's expected utility, we need both the utility function for all possible resulting states of the world and the subjective probability function for the occurrence of those states for each available course of action (Savage, 1972). The expected utility of each course of action is calculated by multiplying the utility of each outcome by its probability of occurring under that course of action. The actor's choices will be consistent with the decision rule of choosing the course of action with the greatest expected utility. To characterize an expected utility model, it is necessary to define the possible states of the world, the choices available to the actor, the utility function for each possible state, and the (subjective) probability distribution for the occurrence of each state of the world for each available choice.

In this model, the possible states of the world correspond to the different levels of policy concessions that could result from a war. It is assumed that nations fight wars to change the policies of their opponents whether those policies be territorial, political, or economic in nature. This range is bounded on both ends because there are physical limits on the policy concessions that can be imposed upon an opponent. For convenience, the limit of the concessions the defender can extract from the attacker will be denoted as a and the upper limit of the concessions the attacker can impose as b, with $a < 0 < b$.[1] Zero is chosen to represent the status quo. Furthermore, each level of policy concessions is associated with a military outcome that would allow one side to seize those policy concessions without a peace treaty. For example, the Falklands/Malvinas war has ended without a peace treaty between England and Argentina. The final outcome in policy terms, English control of the islands, has been imposed by a particular military outcome, the defeat of the Argentine garrison. As the military outcome becomes more favorable to one side, that side will either seize further concessions from its opponent or surrender fewer concessions to its opponent as a consequence of the war. War is seen as a struggle to impose military reverses upon the other side in the interest of extracting policy concessions to cease the war against the possibility of even greater defeat and greater concessions if war continues.

1. The present model assumes that all concessions are viewed conflictually; gains by one side must be losses to the other (although they need not be zero-sum). Because the model views the concessions to be offered unidimensionally, there are no tradeoffs between issues open to the intitiator and defender, and thus compromise will be harder to reach than it is in the real world. Of course, we could assume that the policy space within which concessions can be made is multidimensional (as it most certainly is), which will eliminate this problem. Using a multidimensional policy space would also allow a distinction between different types of concessions and their ease of seizure through military means. Some types of concessions, such as territory, are easy to secure without the consent of the defender, while others, such as internal political policies, are not. In general, nations resorting to war to obtain the latter form of concession attempt to impose costs upon the defender to compel abandonment of those policies offensive to the initiator. The present model is unidimensional in its policy space simply because unidimensional policy spaces are much easier to work with; eventually this restriction and the objection that accompanies it will be eliminated (see Morgan, 1984a, 1984b; Morrow, 1984a, 1984b).

The choices that each side faces are structured into a sequence of threat and counterthreat. First, the nation considering initiating a conflict by making a threat of war (either implicit or explicit) must decide if threatening war is a more profitable course than attempting noncoercive means to gain one's ends. For simplicity's sake, it will be assumed that choosing not to use force will produce an outcome of zero automatically. With the initial threat, the initiator of the conflict offers the defender a compromise proposal, a set of concessions that it will accept in lieu of going to war. After receiving a threat, the target of the threat, referred to henceforth as the defender, must choose either to accept the compromise offered with the threat or to resist it and offer a countercompromise. The final choice for war or peace rests with the initiator; it can accept the countercompromise of the defender and end the crisis or reject it and launch a war between the two countries. Although actual crises may involve several exchanges of threats and compromise offers between the two embroiled nations, this model simplifies those interchanges to the above schematic (for a similar approach using a discrete calculation, see Petersen, 1982). In each of the three decisions, each party to the dispute must judge which course of action is likely to produce the greatest benefit for it, and these decisions will be modeled by an expected utility calculation.

Now that the choices available to the actors and the possible outcomes of those choices have been specified, a utility function for the actors' preferences over those outcomes is needed. In this model, an extremely general utility function will be assumed to represent the preferences of the actors in order to preserve the generality of the model. Three assumptions are made about the structure of utility functions; first, the utility functions are assumed to be bounded and continuous. Second, greater military success is preferred to less success. In mathematical terms, this assumption is simply the statement that the initiator's utility function, $u_i(x)$, is increasing everywhere and the defender's utility function, $u_j(x)$, is decreasing everywhere.

$$\frac{\partial u_i}{\partial x} > 0 \text{ and } \frac{\partial u_j}{\partial x} < 0 \text{ for all } a < x < b \qquad [1]$$

The third assumption states that actors have consistent attitudes toward risk across the entire spectrum of military outcomes. Actors are assumed to have the same preferences for risky situations regardless of the size of those risks. Then the utility function of each actor is just defined by its risk attitude up to an arbitrary increasing and bounded function. For convenience's sake, risk-neutral actors will be given a risk attitude of zero, risk-acceptant actors a positive risk attitude, and risk-averse actors a negative risk attitude, with preference for risky situations increasing with the actor's risk attitude. Mathematically, the assumptions about the effects of risk attitudes on utility functions are as follows:

$$u_i(x) = u(r_i,x) \text{ and } \frac{\partial u_i}{\partial r_i} < 0 \text{ and } r_i \bullet \frac{\partial^2 u_i}{\partial x^2} > 0 \text{ for } r_i \neq 0 \quad [2]$$

The first statement is just that utility functions are determined by an actor's risk attitude up to an arbitrary increasing and bounded function. Second, utility for a particular outcome drops as risk attitude increases. Finally, risk-acceptant actors have concave upward utility functions and risk-averse actors concave downward utility functions. The last two assumptions imply that risk-neutral actors have linear utility functions with respect to military success. The second and third assumptions combine to produce utility functions similar to those pictured in Figure 1.

Referring to Figure 1, the assumptions about the actors' utility functions can be demonstrated. Figure 1 presents three possible utility functions for nation i. The greater the value of $u_i(x)$, the more strongly nation i values the outcome x. The least favorable outcome for all three possible initiators is point a on the left; the most favorable outcome is b on the right. The actors' utilities for the extreme outcomes are set to the same values to facilitate comparison. Actors are assumed to prefer greater levels of success to lesser levels of success, and consequently all three utility functions increase as the outcome x grows in value. The curvature of each utility function induces each actor's preference for risk. The risk-acceptant actor is willing to take risks the other actors would not, and thus, the ratio of the utility differences between winning and the status quo ($u_i(b) - u_i(0)$) and the status quo and losing ($u_i(0) - u_i(a)$) is greater for the risk-acceptant actor than the risk-neutral actor, which in turn is greater than that of the risk-averse actor. Each actor's utility for an outcome is determined only by that outcome and its risk attitude.

Other parties than the initiator and defender can have an effect on the outcome of a war. Because third parties to a conflict will base their decision whether or not to intervene in the conflict on their preferences for the

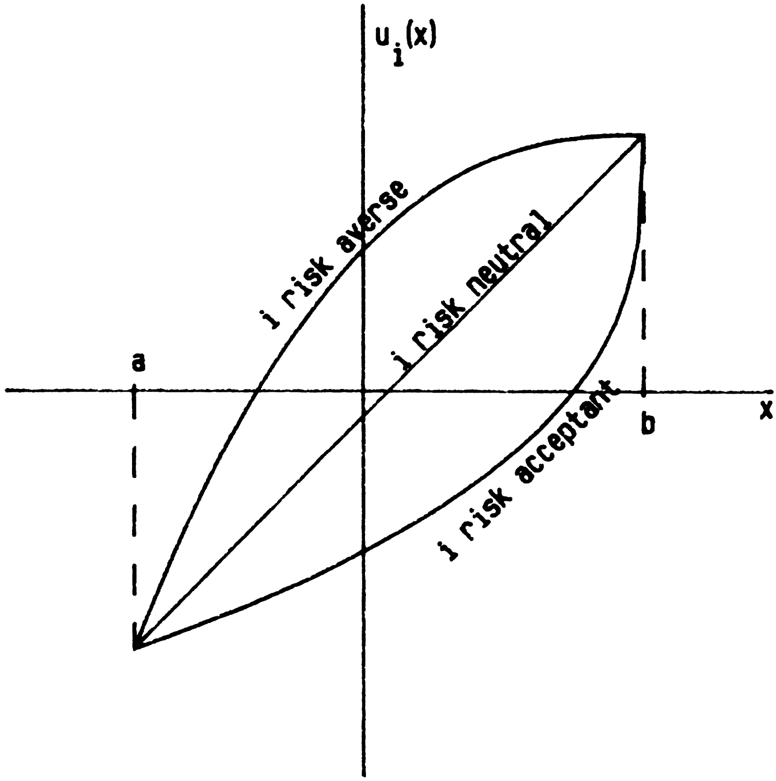

Figure 1

possible outcomes (Altfeld and Bueno do Mesquita, 1979), their utility functions for the different outcomes must also be postulated. Each third party nation k will be assumed to have a preferred outcome, x_k, within the set of possible outcomes from a to b. Around this ideal point, a third party nation k's utility for an outcome will be assumed to depend only on the distance of that outcome from its ideal point and nation k's risk attitude and will decrease as the distance increases. If $d_k = | x - x_k |$, the distance between an outcome and nation k's ideal point, then the assumptions about third party utilities can be summarized as follows:

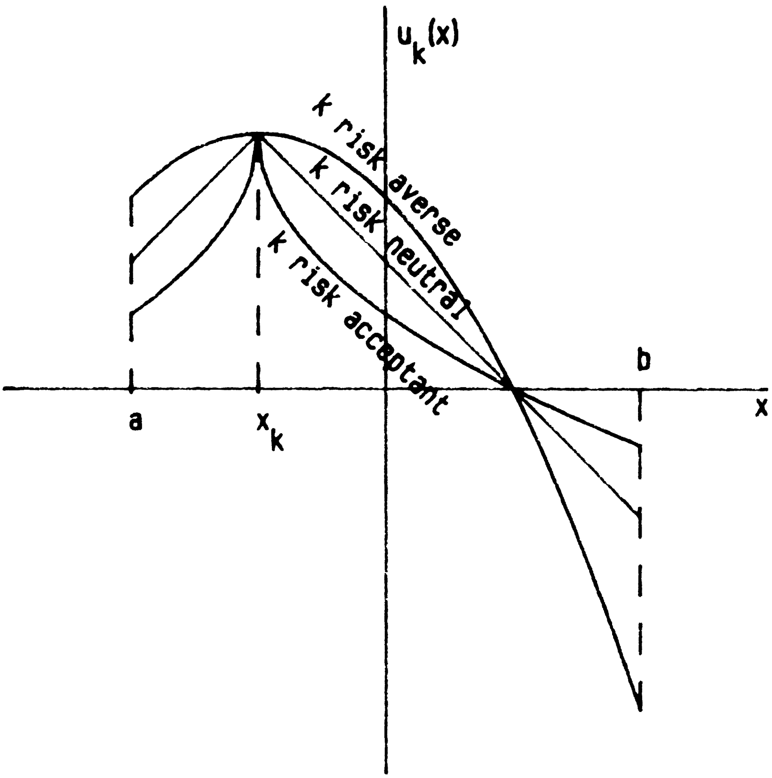

Figure 2

$$\frac{\partial u_k}{\partial d_k} < 0 \quad \text{and} \quad \frac{\partial u_k}{\partial r_k} < 0 \quad \text{and} \quad r_k \frac{\partial^2 u_k}{\partial d_k^2} > 0 \qquad [3]$$

The shapes of third party utility functions according to nation k's risk attitude is similar to those for the initiator and defender; sample utility functions are given in Figure 2.

Figure 2 illustrates three possible utility functions for a third party with ideal point x_k. The utility of x_k is the greatest because of its preferred status. The actor's utility for an outcome drops off as the distance between that outcome and x_k grows, and the utility drop-off is symmetrical around nation k's ideal point. Once again, the willingness of nation k to take risks is reflected in the curvature of its utility function.

A comparison between the utility functions assumed here and Bueno de Mesquita's assumptions about national preferences in *The War Trap* (1981a) will illustrate the connection between the two. Bueno de Mesquita's theory assumes that only two outcomes—victory and defeat—are possible, and as a result of that assumption he needs utilities for only those two outcomes. Equating his utilities of victory and defeat with the utilities for maximal victories and defeats in the continuous model, we have the following:

$$u_i(b) = -u_i(a) = u_{ii} - u_{ij} + \Delta(u_{ii} - u_{ij}) \qquad [4]$$

$$u_j(a) = -u_j(b) = u_{jj} - u_{ji} + \Delta(u_{jj} - u_{ji})$$

Furthermore, because the utility functions in the continuous model are determined only up to a linear transformation, they can be transformed to fit the above equality. However, unlike the model presented here, Bueno de Mesquita does not enter the risk attitudes of the involved nations in his utility estimates as curvatures in the utility functions, as is typical in expected utility analysis (as pointed out by Majeski and Sylvan, 1984, in their critique of *The War Trap*), prefering to reflect the reactions of actors to risk and uncertainty in four different ad hoc decision rules. As a result, the continuous model provides not only greater descriptive ability but also a more complete logic in its treatment of risk than the discrete model Bueno de Mesquita employs.

Bueno de Mesquita corrects this problem by introducing curved utility functions in a subsequent article (Bueno de Mesquita, forthcoming). However, he still assumes that the status quo (zero) will be the midpoint between the limits of victory (b for the initiator) and defeat (a for the initiator), unless the parties expect a change in the other's policies in the absence of war. There is no reason to assume this; there are other reasons why the status quo might be closer to total victory for one side than the other even when the actors expect no change in each other's policies in the near future. For instance, the status quo between the United States and Japan before Pearl Harbor was almost certainly closer to a total American victory than to a total Japanese victory in terms of the disposition of the issues. By ignoring this possible difference, even the improved version of Bueno de Mesquita's expected utility model cannot capture the notion of a status quo power or a challenger that appears in numerous theories of international relations.

To complete the structure of the model, the probabilities of different outcomes for different strategies need to be given. It has already been assumed that if war does not result, either the status quo or a compromise to which both sides agree without fighting will hold with probability one. Because only the expected utility of war has to be calculated, we need only a probability distribution of military outcomes. The resolution of combat will depend only on the military capabilities of the participants. However, the number of participants is unclear before the resolution of a crisis; it is possible that third parties may join the war and have a decisive effect on its outcome. The probabilities of different military outcomes are given by a probability density function that depends on the capabilities of the two conflicting nations and the capabilities and utility functions of all third party nations.

$$P_{ij}(x) = p(cap_i,\ cap_j,\ u_k(x),\ cap_k) \quad \text{for all k} \qquad [5]$$

In a discrete expected utility model, only the probabilities of winning and losing would have to be considered; in this model, the probability of all ranges of outcomes must be considered. The probability of a range of outcomes can be calculated from a continuous probability distribution by integrating the above density function across that range of outcomes; these probabilities can then be treated like the probabilities of discrete outcomes.

The actual probability density is not specified in order to preserve the generality of the model. It is entirely possible that the p function that determines the probabilities of different outcomes would change across different time periods and areas of the globe. It might be that during periods when defense is dominant in warfare, the p function would be less responsive and cluster outcomes around zero. It is also possible that the p function responds to capabilities in a nonlinear fashion; for example, the outcome of a war between two weak nations might be more decisive than between two powerful nations with an equal ratio of capabilities because larger military establishments create a situation where neither side can prevail quickly and decisively. In any case, it is not my intention to pursue these points here; instead, I merely wish to emphasize the advantages to defining the probability distribution of outcomes in only the most general terms. This approach allows conclusions to be drawn from the general model that will hold under many different operationalizations of the variables.

With that brief digression aside, let us return to the assumptions that will be made about the probabilities of the outcomes. The total probability of all possible outcomes is one, or formally,

$$\int_a^b p_{ij}(x)\ dx = 1 \qquad [6]$$

The expectation of the density function is useful in characterizing the changes in the distribution of outcomes. The expectation is the magnitude of each outcome times the probability of that outcome occurring summed across all possible outcomes, which in this case is just the following integral:

$$E(p_{ij}) = \int_a^b x\ p_{\cdot ij}(x)\ dx \qquad [7]$$

More powerful nations will be more likely to prevail in combat; so as a nation's capabilities increase, the expectation will shift in its favor.

$$\frac{\partial E(p_{ij})}{\partial cap_i} > 0 \quad \text{and} \quad \frac{\partial E(p_{ij})}{\partial cap_j} < 0 \qquad [8]$$

This simple and reasonable assumption describes the effect of the capabilities of the initiator and defender on the probability density. The effects of third parties are more complicated because they depend not only on the capabilities of the particular nation but also on its preferences for the outcomes as expressed in its utility function. Briefly, a third party's effect on the probability of an outcome will increase as its capabilities increase and as the strength of its preference for intervening increases. To define the effects of a single third party on the probabilities of different outcomes, an additional function, $p_{ij}^{(k)}$, the probability distribution given k neutral, is needed. This function is equal to p_{ij} with the capabilities of nation k set equal to zero and gives the probabilities of outcomes if nation k does not involve itself in the war. By examining the difference between these two functions, the effects of nation k's participation on the war can be characterized in terms of which outcomes it makes more likely and which outcomes it makes less likely. First, as the military capabilities of nation k increase, the magnitude of its effect will increase, or, in mathematical terms,

$$\frac{\partial |p_{ij} - p_{ij}^{(k)}|}{\partial \, cap_k} > 0 \qquad [9]$$

The absolute value of the difference of the two probability densities gives the magnitude of nation k's effect on the outcome of the war; the larger this difference, the more effect nation k's participation will have. More powerful third parties will have a larger effect on the outcome of a war, all else being equal, than less powerful third parties. Second, nation k's decision whether or not to involve itself in the war and how deeply to commit itself to the fighting will hinge on whether its involvement will produce an outcome more to its liking (Altfeld and Bueno de Mesquita, 1979). The magnitude of a third party's effect on the probability of a particular outcome will increase as the difference between its utility for that outcome and its expected utility for the war if it does not intervene increases. This principle is rather messy to put in mathematical terms; instead the effects of the determinants of nation k's utility on the probabilities of outcomes will be deduced. First, as the distance between an outcome and k's preferred outcome increases, the difference between the two probability densities will decrease.

$$\frac{\partial (p_{ij} - p_{ij}^{(k)})}{\partial d_k} < 0 \qquad [10]$$

The effect of a third party k will be to make the outcomes around k's ideal point more likely; it will assist the side that is attempting to move the final outcome closer to its ideal point until further assistance would allow that side to move the outcome past its ideal point, where k will suspend its aid or assist the other side in an effort to produce its desired outcome. Second, the effects of k's risk attitude on its involvement vary. In general, as k's risk attitude increases, k will commit more resources to the struggle and consequently have a more significant impact on the outcome, but in turn, k's participation will be concentrated to achieve an outcome very close to its ideal point. Risk-acceptant third parties will have a greater expected effect on the outcome, but they will concentrate their efforts to achieve results closer to their ideal point than equally powerful risk-averse third parties. This concludes the description of the assumptions of the model; now is the time to discuss the implications of it.

CONCLUSIONS ABOUT THE INITIATOR AND DEFENDER

According to expected utility theory, a course of action will be chosen only if that course produces a greater expected utility that any other course available. In this model, the initiator is confronted with a choice at the start of a crisis: whether or not to threaten the defender with war. The option of war will be chosen only if the initiator's expected utility for war exceeds that of the status quo.

$$E(U_i)_j > u_i(0)$$

$$\int_a^b u_i(x)\ p_{ij}(x)\ dx > u_i(0) \qquad [11]$$

To find nation i's utility for war, the probability of each outcome is multiplied by i's utility for that outcome and those products are then summed across all possible outcomes. For a continuous probability distribution, the summation is performed by integration across the possible outcomes. If the above condition is not met, then i will not consider war.

We would like to be able to simplify the above condition to draw out its implications. There exists a military outcome, call it x_{i1}, such that nation i will be indifferent between x_{i1} and going to war (see appendix).

$$\int_a^b u_i(x)\ p_{ij}(x)\ dx = u_i(x_{i1}) \qquad [12]$$

The point x_{i1} represents i's indifference point between war and compromise; if j makes concessions to i comparable to those that i could extract from j after achieving a military success of the level of x_{i1}, then i would be indifferent between fighting and accepting those concessions. Obviously, any level of concessions by j greater than x_{i1} would lead i to accept those concessions instead of going to war, and any level less than x_{i1} would lead i to reject those concessions and choose to launch a war instead. Given the choice of initiating a war, nation i will only consider it if $x_{i1} > 0$. Put simply, i will only initiate a crisis if it expects to gain concessions beyond the status quo from the conflict.

Turning to the defender's perspective, it also has a military outcome equivalent to its expected utility for war, call it x_{j1}.

$$\int_a^b u_j(x)\ p_{ij}(x)\ dx = u_j(x_{j1}) \qquad [13]$$

Any proposed level of concessions less than x_{j1} will be acceptable to j in lieu of going to war. Then agreement short of war will be possible only if there exists a point x_o such that $x_{i1} < x_o < x_{j1}$, which of course implies that $x_{i1} < x_{j1}$. In layperson's terms, there must exist a compromise that is preferable to fighting to both sides.[2] In addition, the defender must make concessions to the initiator because x_{i1} must be greater than zero for otherwise the initiator would not have initiated the conflict because the status quo would have been preferred to war.[3]

This conclusion is hardly surprising; for war to occur both sides must prefer war to any compromise offered by their opponent. The level of acceptable concessions provides a tool for determining which side is likely to prevail in a crisis whether that crisis escalates to war or not. The larger the difference an actor perceives between the status quo and war, the more aggressively that actor will pursue coercive strategies in crises, or, in more common terms, the more resolve the actor will possess. In terms of the model, as x_{i1} increases, nation i will possess greater resolve to pursue the crisis through an advantageous conclusion and, consequently, will be more likely to prevail in that crisis. Maoz (1983) presents evidence that supports this position by demonstrating that those nations that expressed the greatest willingness to wage war through the hostility of their actions prevailed in crisis situations. Because the level of acceptable concessions represents an expectation of military success, the level of expected military success will rise with increases in the level of acceptable

2. This construction is essentially identical to Wittman's (1979), particularly if it is applied to negotiations during a war. He then uses the level of acceptable concessions to analyze under what conditions agreements to end wars can be reached. However, instead of explicitly defining a nation's expected utility for war and proceeding, he makes arguments about how these expected utilities will change with changes in strategies during war. As a result, all of his conclusions about the effects of different strategies on the final outcome of a war also follow from this model.

3. If we include bluffing strategies in the model, then it is possible that a crisis could be initiated by a nation that did not expect to win a war if it believed that its target was very likely to make the demanded concessions. When a nation makes a bluff threat, it attempts to use the defender's fear of losing a war to gain a quick concession short of war. However, if the defender chooses to stand firm, the initiator of the threat will probably back down. To model this possibility, the initiator's expected utility would include a term for expected concessions from the defender short of war.

concessions. So success in crises should correspond to holding higher levels of acceptable concessions.

Under what conditions will the initiator's level of acceptable concessions increase? First, as the correlation of forces becomes more favorable to the initiator, its demands will increase. As the expectation of military outcome increases, the initiator's expected utility must increase, which will lead to a higher level of concessions being demanded. The expectation of military outcome will increase as the initiator gains in capabilities relative to the defender and as the initiator gains powerful allies who will assist it in gaining its ends through military means.

Second, as the initiator's risk attitude increases, it will demand higher levels of concessions (see appendix for a proof of this statement). Although the initiator's expected utility for war will drop as its risk attitude increases, this decrease will be more than offset by a greater drop in the initiator's utility for the status quo. Risk-acceptant actors will demand higher levels of concessions because they are willing to risk defeat or stalemate in order to gain a chance at victory which they value much more highly than either defeat or stalemate. In Figure 3, the area under each of the utility functions gives the expected utility for war of a risk-averse and risk-acceptant actor (assuming a uniform distribution of outcomes); note that the area is larger under the curve of the risk averse actor. Then for each actor, x_{i1} is the outcome whose utility is equal to the area under the curve. Each actor's value for the status quo (0) can also be read off Figure 3, and the risk-acceptant actor has a much lower value for the status quo than the risk-averse actor's demands because its utility for the status quo is lower, even though its expected utility for war is also lower (recall that all utility evaluations are relative).

Third, as the position of the status quo deteriorates relative to the initiator's position, the level of acceptable concessions will rise relative to the status quo, and the initiator's resolve will increase. Holding the distance between the extreme outcomes a and b constant, moving the status quo toward a will reflect the effect of a status quo that is less favorable to nation i. If we hold the distance between the extreme outcomes, nation i's utility function, and the probability distribution of outcomes constant, then x_{i1} will not shift relative to the extreme outcomes, a and b. Then as shown in Figure 4, when the status quo outcome is moved closer to a, the difference between it and x_{i1} will increase, simultaneously increasing the resolve of the initiator. There are three sources of resolve then: one, a more favorable probability distribution; two, a greater willingness to ac-

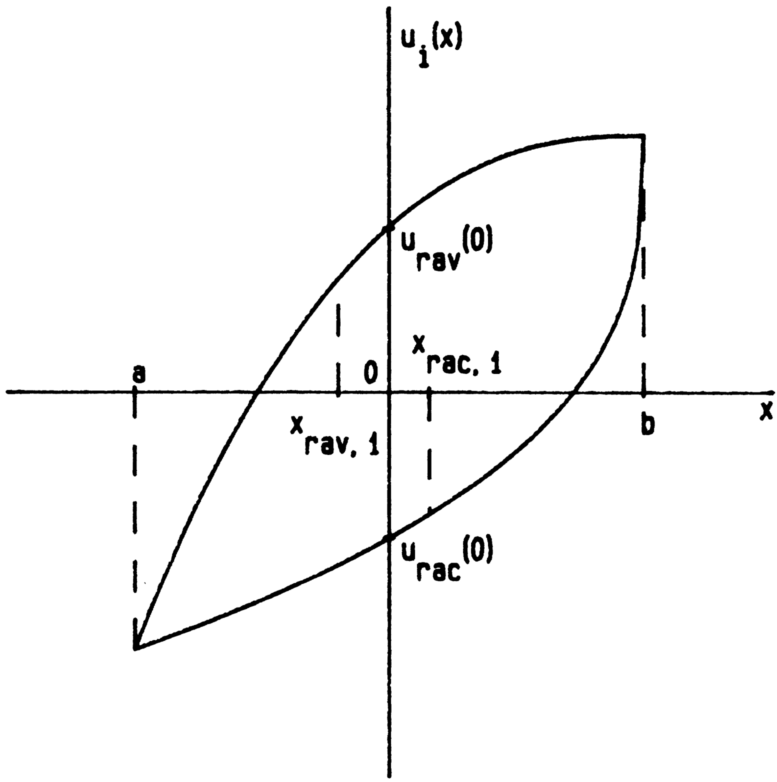

Figure 3

cept risks; and three, a less favorable disposition of the status quo. War then is a gamble over the possible outcomes; like gambles, war becomes more attractive as the likelihood of winning increases, as one's value for winning increases, and as the alternative to fighting becomes less attractive.

The level of acceptable concessions also provides an additional insight about the likelihood of the initiation of wars. As Wittman (1979) points out in his article on the termination of wars, the distribution of power is irrelevant to the termination or initiation of wars. The distribution of power between two actors determines the probability distribution of outcomes that the actors face when considering war. Assuming that both ac-

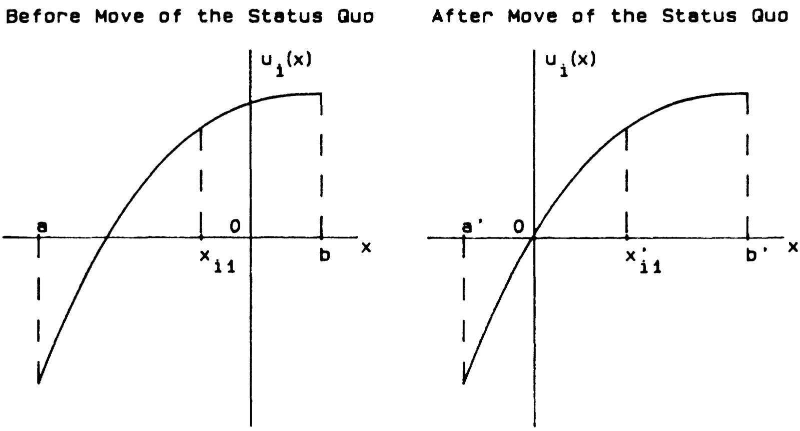

Figure 4

tors perceive the same probability distribution, changes in that probability distribution toward one of the actors will change both actors' acceptable level of concessions in the same direction. For example, if the distribution of power shifts toward nation i, making outcomes closer to b more likely, both x_{i1} and x_{j1} will increase. Whether or not war becomes more or less likely after this shift depends upon the exact magnitude of those shifts, and those magnitudes depend upon both actors' risk attitudes as well as the shift in the probability distribution. Thus it cannot be said that either a balance or a preponderance of power leads to war without knowledge of the actors' risk attitudes.[4]

These calculations agree with the arguments that Bueno de Mesquita presents on escalation to violence in crises. He argued that the likelihood of a crisis escalating to war could be deduced by comparing the expected utilities of the conflicting nations. If the absolute value of the defender's negative expected utility was greater than the attacker's expected utility, he argued that both sides would be better off negotiating than fighting, and thus war would not occur (Bueno de Mesquita, 1981a). He then went on to compare the defender's and attacker's expected utilities in wars to estimate the cost of the war in lives. However, as Zagare (1982) pointed

4. Bueno de Mesquita (1981b) demonstrates the same point using a computer simulation that varies both the distribution of power and risk attitudes.

out in his review of *The War Trap*, both of these calculations entail interpersonal comparisons of utility. The theory presented in this article allows the same conclusions while avoiding the direct comparisons of utilities. Here the comparisons are made in terms of the concessions each side will demand, instead of the expected utility each side holds for war as a strategy. Because all comparisons across nations are made in terms of acceptable levels of concessions, the theory compares intensities of preferences through actions rather than utilities, eliminating the problem of interpersonal comparisons of utilities.

Once a war starts, the model can provide us with further information about the likely outcome of the fighting. There exists another point, x_{i2}, generally different from x_{i1} with the following property:

$$\int_a^b u_i(x)\ p_{ij}(x)\ dx = u_i(a) \int_a^{x_{i2}} p_{ij}(x)\ dx + u_i(b) \int_{x_{i2}}^{b} p_{ij}(x)\ dx \qquad [14]$$
$$= u_i(a)\ (1 - p_i) + u_i(b)\ (p_i)$$
$$\text{where } p_i = \int_{x_{i2}}^{b} p_{ij}(x)\ dx$$

The point x_{i2} allows a reduction of the continuous probability model to a discrete probability model over the outcomes of victory and defeat where $u_i(b)$ is i's utility for victory and $u_i(a)$ its utility for defeat. Then x_{i2} represents the level of policy concessions that i will need to consider the war victorious (or just its level of acceptable victory). The integral of p_{ij} from x_{i2} to b is just the probability that i will consider the war victorious and is set equal to p_i, i's probability of victory. There is a corresponding point for the defender, x_{j2}, which allows the above reduction of its utility calculus to a discrete calculation.

This reduction of a continuous-outcome to a discrete-outcome utility calculus provides the key to estimating expected utility calculations. Every expected utility is equivalent to an expected utility calculation over the outcomes of winning and losing and, furthermore, can be empirically estimated by that discrete-outcome expected utility. Additionally, the utilities of winning and losing can be chosen arbitrarily, and therefore, an actor's expected utility for war depends only on its subjective probability of winning, p_i. This subjective probability of winning must include the warring parties' military capabilities, possible third parties' military capabilities, those possible third parties' preferences for the warring parties, and the risk attitude of the deciding actor. Estimating an actor's subjective probability of winning a war is identical to estimating expected utility for that actor of the war because the utilities attached to the outcomes of winning and losing are arbitrary.[5]

The implications of this conclusion for empirical tests of expected utility theories are important. For instance, Wagner (1984) argues that Bueno de Mesquita (1981a) does not test an expected utility theory in *The War Trap* but rather a balance of power theory, because Bueno de Mesquita's empirical calculation of a nation's expected utility reduces to an estimate of whether the nation and its allies, discounted for each ally's reliability, possess greater capabilities than its opponent and its allies (also discounted for their reliability). But such a calculation is just an estimation of that nation's probability of victory, and thus is equivalent to an expected utility calculation. However, that calculation omits the effects of the nation's risk attitude on its subjective probability of winning, and consequently cannot capture that subjective probability completely. The addition of risk attitudes in a subsequent article by Bueno de Mesquita (forthcoming) rectifies this problem. (Nevertheless, it should be pointed out that Bueno de Mesquita enters the actors' risk attitudes through their utility functions, while the above results suggest that those risk attitudes should be used to alter the actors' subjective probabilities of winning and losing.)

Beyond allowing a reduction to a discrete utility calculus that is far easier to calculate and manipulate, these levels of subjective victory lead us to a new insight. Because they are not generally equal to the level of concessions demanded prior to the start of war, they explain why some countries increase their demands after the start of a war while others reduce theirs. If $x_{i1} < x_{i2}$ for some nation i, then i will increase its demands for concessions after hostilities commence. What were acceptable concessions to nation i before the fighting started no longer are. Once hostilities begin, some nations will escalate their demands because they believe greater demands are necessary to consider the war victorious, while other nations will attempt to cut their losses and accept less favorable outcomes than they would have before the war started. We would then like to know which cases will head in which direction once war commences; the answer depends on the risk attitude of the participants and the probability distribution of outcomes. In general, however, risk-acceptant actors will increase their demands while risk-averse actors will reduce theirs. In many cases then, a nation could gain more favorable

5. However, calculating whether or not an actor prefers war to the status quo still hinges upon its utility for the status quo, which is a function of its utilities for winning and losing and the actor's risk attitude.

outcomes by pushing risk-averse opponents toward war while appeasing risk-acceptant opponents.

An excellent example of this phenomenon of shifting levels of acceptable victory is the American intervention in the Vietnam war. The United States demonstrated its military superiority over North Vietnamese forces on the battlefield (Palmer, 1978; Summers, 1981). However, American military successes on the battlefield were not translated into a favorable settlement to the dispute; in fact, the Paris peace settlement contained a number of provisions that were distinctly unfavorable to American interests in Vietnam. In terms of the theory, if the United States is considered to be risk averse, its level of acceptable concessions to conclude the war would be lower than its level of acceptable concessions before American troops were sent into combat. Under this explanation, once the North Vietnamese were able to prevent a quick American military victory, the American leadership was willing to consider negotiated outcomes that were objectively unfavorable but satisfied its level of acceptable victory. Historically, President Nixon even proclaimed the peace settlement as "Peace with Honor," underlining its acceptability to both the American leadership and populace.[6]

CONCLUSIONS ABOUT THIRD PARTIES: THE MULTILATERAL DECOMPOSITION

The conclusions already drawn hold for the calculations of the defender and the initiator regardless of the participation of third parties. The effects of third parties have already been built into the probability distribution of the outcomes, and their intervention is assumed to have no effect on the participants' values for the outcomes of the fighting. In the continuous model, the only value of a third party is to make desirable outcomes more likely than undesirable ones for an actor assessing the

6. The multidimensional model provides a more through explanation of the Vietnam war. The goal of American policy in Vietnam was to stop the North Vietnamese infiltration of troops and supplies into South Vietnam aimed at overthrowing the government there. Instead of choosing to stop this infiltration through direct military means, the United States bombed North Vietnam to compel them to desist their infiltration and fought on the ground in South Vietnam to support its government. When the United States found that North Vietnam was willing to accept very high costs rather than cease its infiltration into the South, the United States searched for a compromise to end its participation in the war. Being risk averse, the United States was willing to accept a final settlement that allowed North Vietnam to continue its infilitration into South Vietnam.

third party's expected utility. Now we would like to examine how initiators and defenders assess the effects of third parties on the conflicts they face. In *The War Trap*, the marginal contribution of each third party in utility terms is assessed by multiplying each actor's net utility for the participation of that third party ($u_{iki} - u_{ikj}$) by the relative capability of that nation to the combatants ($p_{ik} + p_{jk} - 1$). Wagner (1984) argues that the proper way to calculate the expected utility with third parties is to decompose the probability of winning and losing into a series of conditional probabilities of winning and losing for each possible set of third parties that could intervene multiplied by the probability that those third parties would intervene. These effects are already subsumed into the probability distribution of outcomes in the continuous model. The conditions under which Bueno de Mesquita's approach and the continuous model are the same will be deduced to check the logical consistency of the former.

First consider a case in which only one third party nation exists. The probability of different outcomes depends only on the capabilities of all three combatants and the utility function of the third nation. As in the previous discussion of the motivations of third parties, we consider the function $p_{ij}^{(k)}$, which is equal to p_{ij} with the capabilities of the third party k set to zero. This function, $p_{ij}^{(k)}$, corresponds to the distribution of results that will occur if nation k is either unable to influence events or indifferent to them, or in other words, if the resulting war is bilateral with absolutely no outside intervention. Returning to the initiator's expected utility for war considering all possibilities,

$$
\begin{aligned}
E(U_i)_j &= \int_a^b u_i(x)\ p_{ij}(x)\ dx \qquad [15] \\
&= \int_a^b u_i(x)\ p_{ij}^{(k)}(x)\ dx + \int_a^b u_i(x)\left(p_{ij}(x) - p_{ij}^{(k)}(x)\right) dx \\
&= u_i(b)\ (p_{ij}^{*}) + u_i(a)\ (1 - p_{ij})^{*} + \int_a^b u_i(x)\left(p_{ij}(x) - p_{ij}^{(k)}(x)\right) dx
\end{aligned}
$$

The initiator's calculation then is a discrete, bilateral calculation of expected utility against the defender; where $u_i(b)$ is the initiator's utility for winning, $u_i(a)$ is the initiator's utility for losing, and p^{*}_{ij} the probability that the initiator will win a bilateral war with the defender; plus an adjustment for the expected effects of the third party k (the integral in the final equation). Furthermore, third parties fall into two classes, those that favor the initiator and those that favor the defender, and for each third party k, there exists a point x_{ik} such that (see the appendix for the details of this derivation)

$$\int_a^b u_i(x)\left(p_{ij}(x) - p_{ij}^{(k)}(x)\right) dx = \left(\int_{x_{ik}}^b p_{ij}(x) - p_{ij}^{(k)}(x)\, dx\right)(u_i(b) - u_i(a)) \quad [16]$$

$$= p_{ik}(u_i(b) - u_i(a)) \qquad \text{if k favors i}$$

or, if k favors j,

$$= \left(\int_a^{x_{ik}} p_{ij}(x) - p_{ij}^{(k)}(x)\, dx\right)(u_i(a) - u_i(b))$$

$$= p_{ik}(u_i(a) - u_i(b)) = -p_{ik}(u_i(b) - u_i(a))$$

The effect of any third party is quite simple then; it is just the probability that k plays a decisive role in the war (p_{ik})—that is, turns a defeat for the side it favors into a victory—multiplied by the difference in utility that change produces for the initiator.

The point x_{ik} gives the initiator's expected settlement if the third party should intervene in the war. The expected settlement after intervention incorporates the moderating effect of the third party on the demands of the victorious side. The initiator will consider any outcome greater than this level as an advantage achieved by the intervention of the third party into the war. The farther x_{ik} is from the ideal point of the side it favors (a if the defender, b if the initiator), the greater k's effect on the outcome (p_{ik}) will be. The defender's calculation can also be broken down in an identical fashion, so both the defender and initiator will consider a third party as either a friend or an enemy and assess its effect on the conflict in terms of how likely that country's intervention is to reverse the tide of battle.[7]

When will a third party have a greater effect on the outcome of a war? First, the more powerful a nation is, the greater its effect on the outcome will be, all else equal. This conclusion follows from equation 9. Second, because k's contribution depends on the strength of its preference for the outcomes it can produce by its intervention, k's participation will increase as the strength of its preference over the outcome increases. Although this statement is obvious, it has a number of nonobvious corollaries. Allies with preferences very similar to the initiator or defender are very desirable for them because those allies will be likely to use their capabilities to help their preferred combatant. Strong allies with moderate preferences can be as much of a liability as an asset to a risk-acceptant combatant because they may deny the combatant the extreme military victory that is strongly desired. Risk-averse allies, on the other hand, are generally quite desirable provided that they intervene; they are generally reluctant to commit themselves unless defeat is likely without

7. It is entirely possible that both the initiator and the defender could consider a particular third party as a friend (or as an enemy). If a third party is very close to neutrality between the two nations, it is quite probable that this would occur.

aid, but if they do turn the tide of combat, they are likely to provide their once-embattled ally with the liberty of setting its own terms.

At first glance, equation 16 looks quite different from Bueno de Mesquita's (1981a) calculation for the expected utility derived from third parties. Setting his expression for the expected utility derived from a third party to equation 16, we have the following:

$$\begin{aligned}(p_{ik} + p_{jk} - 1)(u_{iki} - u_{ikj}) &= (u_i(b) - u_i(a))(p_{ik}) \\ &= ((u_{ii} - u_{ij}) - (u_{ij} - u_{ii}))(p_{k\ \text{intervenes}})(cap_k) \\ &= 2(u_{ii} - u_{ij})(p_{k\ \text{intervenes}}(cap_k)\end{aligned} \quad [17]$$

where u_{ii} is i's utility for its own policies, u_{ij} is i's utility for j's policies, p_{ik} is the probability that i wins given that k aids i, p_{jk} is the probability that j wins given that k aids j, u_{iki} is i's perception of its utility to be derived from k, u_{ikj} is i's perception of the utility j will derive from k's participation, and cap_k is just k's capabilities relative to i and j. Setting cap_k equal to $p_{ik} + p_{jk} - 1$ (as Bueno de Mesquita does), the probability that k intervenes can be found.

$$P_{k\ \text{intervenes}} = \frac{u_{iki} - u_{ikj}}{2(u_{ii} - u_{ij})} \quad [18]$$

Bueno de Mesquita's theory will be logically consistent with the more general continuous expected utility model only if the above relationship is true. Personally, I find this statement quite plausible and acceptable: nevertheless, it is an unstated assumption in the framework of the expected utility theory as formulated by Bueno de Mesquita, and it must be justified. Empirical justification could be provided by estimating the probability that a given third party intervened (as was done in Altfeld and Bueno de Mesquita, 1979) and comparing those estimates to the estimates obtained using equation 18. In the absence of such empirical justification, some argument must be provided to support this implicit assumption within the theory. Alternatively, the multilateral portion of the expected utility could be calculated using the estimated probability of intervention.

Beyond just knowing the effects of allies on the calculations of the initiator and defender, we would like to know how and when the complete calculation can be decomposed into a discrete outcome calculus. It has already been shown that this reduction is not only possible but positively simple for the entire calculation. Now we would like to decompose the calculation to expose the marginal effect of each third party individually as Bueno de Mesquita's (1981a) calculation does. To do this, we need the initiator's probability of victory in a bilateral war. Setting the capabilities of every third party to zero, we produce a probability distribution for bilateral war, $p^{(0)}_{ij}$. The resulting bilateral expected utility calculated using this distribution can be reduced to a discrete expected utility calculation as already shown by finding the points x_{i2} and x_{j2} for that calculation. Furthermore, the effect of each third party on this calculation can be found by the discrete calculation given in equation 16. These calculations can be summed to find the expected marginal effect of a set of third parties only if the effect of each third party on the probability distribution is additive and cumulative. Hence, a decomposition of a continuous expected utility calculation into a discrete calculus where the effect of each third party is separately calculated can only be performed if the assumption of independence of intervention and effect holds for all third parties; that is, there are no synergistic effects of multiple interventions on the outcome of the war. For instance, the effect of Britain and France together on the outcome of a war must be equal to the sum of their separate effects. Furthermore, each country's decision to intervene must be independent of any other country's decision to intervene. In general, neither of these assumption is likely to be true. For instance, the effects of each country that intervened in the Austro-Serbian war, which mushroomed into World War I, were clearly not linear because the addition of many major powers on both sides almost certainly drove the outcome toward the military stalemate that resulted. Furthermore, the intervention of each successive country into World War I was certainly not independent of the other interventions. Nevertheless, in the absence of a more detailed examination of interventions into war and the impacts of third parties on the outcomes of wars, these simplifying assumptions are acceptable at this point.

To recap the argument, nations calculate their expected utility for war and compare it to their utility for other strategies when deciding whether or not to go to war. Although an expected utility for war must be calculated over the continuous distribution of possible outcomes, there exists an equivalent expected utility calculation over the discrete outcomes of winning and losing that separates the effects of third parties into the individual effects of each on the probability of the initiator winning.

$$E(U_i)_j = \int_a u_i(x)\, p_{ij}(x)\, dx = u_i(a)\,(1 - p_i^{(0)}) + u_i(b)\,(p_i^{(0)}) + \sum_{k \text{ favors } i} p_{ik}\,(u_i(b) - u_i(a)) \qquad [19]$$

$$- \sum_{k \text{ favors } j} p_{ik} (u_i (b) - u_i (a))$$

In equation 19, which gives only the initiator's expected utility (the defender's is similar), $p_i^{(0)}$ is the initiator's subjective probability of victory in a bilateral war, $u_i(b)$ the utility of victory, $u_i(a)$ the utility of defeat, and p_{ik} the probability that third party k has a decisive effect on the outcome (which is specific to each third party). Because the discrete equation only requires the estimation of discrete probabilities and utilities rather than entire utility functions and complete probability distributions, it is probably easier to operationalize than the continuous model. Nevertheless, those estimates used in the equivalent discrete expected utility depend upon the underlying continuous probability distribution, and so any estimates used in operationalizing a discrete model implicitly depend upon assumptions about the underlying continuous probability distribution. In the absence of empirical evidence suggesting one set of distributions over another, assumptions on the distributions of outcomes should be made explicit in order that they can be examined for their credibility. Different assumptions will lead to different operationalizations of the probabilities of victory and defeat in the estimated discrete equation, and it is important that the assumptions underlying any one operationalization be know.[8]

CONCLUSION

By deducing the conditions under which a continuous-outcome expected utility model can be reduced to a discrete-outcome expected utility model, several unstated assumptions within the discrete-outcome model have been unearthed. The expected utility calculus presented in *The War Trap* assumes that the probability of intervention in wars by third parties is given by equation 18, the cumulative effect of several interventions on the outcome of a war is just the sum of their separate effects (independence of effect), and the probabilities that third parties intervene are

8. It should be noted that Bueno de Mesquita's operationalization of the discrete probabilities cannot be derived from the reduction of any simple (linear or quadratic) continuous probability distribution, even if the actors are assumed to be risk neutral. This result implies that his probabilities are incorrect for some range of the ratio of capabilities between the warring parties (when the ratio is large, Bueno de Mesquita probably does not give the smaller nation a sufficient probability of winning).

independent of each other (independence of intervention). The operationalization of the probability terms depends upon the probability distribution of military outcomes, and the relationship between the two is presently unclear. The continuous-outcome model also presents us with two conclusions not obvious from the discrete-outcome model. First, risk attitudes are a key determinant in who prevails and by how much in a crisis, whether or not a crisis escalates to war. Second, decisions to demand greater (or lesser) concessions after hostilities commence than before are consistent with a rational model. Finally, the continuous-outcome model eliminates the problem of interpersonal comparisons of utilities that plagued the original theory and demonstrates that the utilities for the intervention of third parties are consistent with the utilities for bilateral outcomes under the assumed probability of intervention by those third parties.

This examination of the expected utility theory demonstrates the purpose of theoretical articulation. By loosening the assumptions of the original theory, unstated assumptions within the body of that theory and additional implications of the expected utility paradigm have been found. One of those additional implications—the importance of risk attitudes for crisis outcomes—is already empirically supported by the Maoz study, and the other—changing levels of victory before and after hostilities—seems historically credible. Furthermore, the theory suggests improved methods of testing the expected utility theory by pointing to the importance of the probability distribution of military outcomes for the operationalization of the probability terms. Testable statements about the effects of changes in this distribution on the conclusions of the theory should be easily deducible. In a sense, the theory presented here subsumes Bueno de Mesquita's expected utility theory, but this case is not a spectacular example of theoretical subsumption because we expected to include the expected utility theory within the more general theory. Theoretical subsumption is crucial to the development of theories when two apparently incongruous theories are reconciled within one framework. Still, the example presented does illustrate several of the advantages of the articulation of theory.

The results of this article have implications beyond just the expected utility theory and lead to the construction of other theories of international relations. In a sense, the expected utility theory holds a special place in the set of theories of international relations because of its ability to explain consistent national action. Because a utility function can be constructed that fits any consistent pattern of actions of a nation, any theory that assumes that national action is consistent must observe the conclusions of the expected utility theory or else be logically inconsistent. For instance, it is simply not logically credible to argue that either a balance or preponderance of power leads to peace—as do the balance of power (Morgenthau, 1973) and power transition (Organski, 1968; Organski and Kugler, 1980) theories—unless one is willing to argue one of the following: (1) nations' actions are not consistent (i.e., some alternatives are viewed as incomparable to others or else cycles of transitivity exist in national preferences); (2) nations do not prefer greater military success to less military success; or (3) there is an empirical relationship between the distribution of power and the distribution of risk attitudes that makes war more likely under one distribution of power than another.[9] The problem with theories that assert that a particular distribution of power leads to war (or peace) is not that any of the above positions is untenable, but rather that the supporting arguments needed to establish those positions are not advanced. Because the expected utility theory provides an extremely general description of national behavior, it provides a logical building block for the construction of systemic theories of international relations, provided that the assumptions of those theories can be reduced to changes in the probabilities and utilities of nations' expected utility calculations. However, theoretical structures built from the expected utility theory can only be as sound as the underlying theory itself and, therefore, it is incumbent for expected utility theorists to build as general and strong a logical foundation as is possible.

APPENDIX

This appendix contains mathematical proofs of several statements in the text that are not obvious from the assumptions of the model (statements of theorems are drawn form Apostol, 1973, and DeGroot, 1975). First, the existence of x_{i1}, the level of acceptable concessions, must be

9. Alternatively, one could argue that there is a relationship between nations' subjective probabilities of winning and the distribution of such power that with a balance of power both sides will overrate their chance of winning, leading to an increase in their level of acceptable concessions and thus a greater chance of war. However, I assumed that the probability distribution of outcomes was known to both actors in this article, and so the present model cannot be used for a formal discussion of the effects of joint optimism or pessimism on the likelihood of war. Still, see Wittman's (1979) discussion of the effects of subjective probabilities on the likelihood of reaching a settlement.

proved. The probability density function, p_{ij}, is always nonnegative; and consequently the cumulative distribution function, which is the integral of the density function, must be increasing and, therefore, of bounded variation. Furthermore, because the utility functions are continuous and the density functions are of bounded variation, the utility functions are integrable with respect to the cumulative distribution functions. In other words, under the assumptions about the utility and density functions made in the article, every expected utility must exist. Furthermore, by the first mean-value theorem for integrals, there must exist a real number x^*, $a \leqslant x^* \leqslant b$, such that

$$u_i(x^*) = \int_a^b u_i(x)\, p_{ij}(x)\, dx \qquad [A1]$$

Let $x_{i1} = x^*$ for u_i as the integrand and $x_{j1} = x^*$ with u_j as the integrand. This proves the existence of the level of acceptable concessions.

To prove the existence of x_{i2} (and by parallel logic, x_{j2}), recall that u_i is an increasing function and that the cumulative distribution function is continuous. By the second mean-value theorem for integrals, there must exist x', $a \leqslant x' \leqslant b$, such that

$$u_i(b) \int_{x^*}^b p_{ij}(x)\, dx + u_i(a) \int_a^{x^*} p_{ij}(x)\, dx = \int_a^b u_i(x)\, p_{ij}(x)\, dx \qquad [A2]$$

Set $x_{i2} = x'$ in this case for u_i as the integrand and $x_{j2} = x'$ with u_j as the integrand. To establish the existence of x_{ik}, note that

$$\int_a^b p_{ij}(x) - p_{ij}^{(k)}(x)\, dx = \int_a^b p_{ij}(x)\, dx - \int_a^b p_{ij}^{(k)}(x)\, dx = 0 \qquad [A3]$$

and thus for any x_{ik} such that $a \leqslant x_{ik} \leqslant b$,

$$\int_a^{x_{ik}} p_{ij}(x) - p_{ij}^{(k)}(x)\, dx = -\int_{x_{ik}}^b p_{ij}(x) - p_{ij}^{(k)}(x)\, dx \qquad [A4]$$

and apply the second mean value theorem for integrals to arrive at the following:

$$\int_a^b u_i(x)\,(p_{ij}(x) - p_{ij}^{(k)}(x))dx = u_i(b)\left(\int_{x_{ik}}^b p_{ij}(x) - p_{ij}^{(k)}(x)\, dx\right) + \qquad [A5]$$

$$u_i(a)\left(\int_a^{x_{ik}} p_{ij}(x) - p_{ij}^{(k)}(x)\,dx\right) = \left(\int_{x_{ik}}^{b} p_{ij}(x) - p_{ij}^{(k)}(x)\,dx\right)(u_i(b) - u_i(a))$$

Finally, it must be shown that x_{i1} and x_{i2} increase as r_i increases. Taking the latter first, because u_i decreases as r_i increases, $E(U_i)_j$ decreases as r_i increases. With $u_i(a)$ and $u_i(b)$ fixed, p_i must decline along with the decline in $E(U_i)_j$, and hence x_{i2} must increase concurrently. To demonstrate the former contention, differentiate both sides of equation 12 in the text to arrive at the following:

$$\int_a^b \frac{\partial u_i}{\partial r_i} p_{ij}(x)\,dx = \frac{\partial u_i}{\partial r_i}\,\frac{\partial r_i}{\partial x_{i1}} = \frac{\partial u_i}{\partial r_i}\left[\frac{1}{\dfrac{\partial u_{i1}}{\partial r_i}}\right] \qquad [A6]$$

The left side of the above equation is negative because $E(U_i)_j$ decreases as r_i increases; the first term on the right side of the equation is negative by assumption; therefore, the second term on the right side must be positive, and x_{i1} must increase as r_i increases.

REFERENCES

ALTFELD, M. and B. BUENO DE MESQUITA (1979) "Choosing sides in wars." Int. Studies Q. 23: 87-112.

APOSTOL, T. M. (1973) Mathematical Analysis. Reading, MA: Addison-Wesley.

BLAINEY, G. (1973) The Causes of War. New York: Free Press.

BUENO DE MESQUITA, B. (forthcoming) "*The War Trap* revisited: a revised expected utility model." Amer. Pol. Sci. Rev.

———(1981a) The War Trap. New Haven, CT: Yale Univ. Press.

———(1981b) "Risk, power distributions, and the likelihood of war." Int. Studies Q. 25: 541-568.

DeGROOT, M. H. (1975) Probability and Statistics. Reading, MA: Addison-Wesley.

LUCE, R. D. and H. RAIFFA (1957) Games and Decisions. New York: John Wiley.

MAJESKI, S. J. and D. SYLVAN (1984) "Simple choices and complex calculations: a critique of *The War Trap*." J. of Conflict Resolution 28: 316-340.

MAOZ, Z. (1983) "Resolve, capabilities, and the outcomes of interstate disputes, 1816-1976." J. of Conflict Resolution 27: 195-229.

MORGAN, T. C. (1984a) "A spatial model of crisis bargaining." Presented at the 1984 annual meeting of the International Studies Association.

———(1984b) "Power differentials, resolve, and bargaining in international crisis: a spatial model." Presented at the 1984 annual meeting of the American Political Science Association.

MORGENTHAU, H. J. (1973) Politics Among Nations: The Struggle for Power and Peace. New York: Knopf.
MORROW, J. D. (1984a) "A spatial theory of international conflict." Presented at the 1984 annual meeting of the International Studies Association.
———(1984b) "Some tests of a spatial theory of international conflict." Presented at the 1984 meeting of the Midwest Political Science Association.
ORGANSKI, A.F.K. (1968) World Politics. New York: Knopf.
———and J. KUGLER (1980) The War Ledger. Chicago: Univ. of Chicago Press..
PALMER, D. R. (1978) Summons of the Trumpet: U.S.-Vietnam in Perspective. San Rafael, CA: Presidio Press.
PETERSEN, W. J. (1982) "Conflict resolution and the status quo." Presented at the 1982 annual meeting of the International Studies Association.
RIKER, W. H. and P. C. ORDESHOOK (1973) An Introduction to Positive Political Theory. Englewood Cliffs, NJ: Prentice-Hall.
SAVAGE, L. J. (1972) The Foundations of Statistics. New York: Dover.
SNYDER, G. H. and P. DIESING (1977) Conflict Among Nations: Bargaining, Decision-Making, and System Structure in International Crises. Princeton, NJ: Princeton Univ. Press.
SUMMERS, H. G., Jr. (1981) On Strategy: The Vietnam War in Context. Washington, DC: USGPO.
WAGNER, R. H. (1984) "War and expected utility theory." World Politics 36: 407-423.
WITTMAN, D. (1979) "How a war ends: a rational model approach." J. of Conflict Resolution 23: 743-763.
ZAGARE, F. C. (1982) "Review of *The War Trap*." Amer. Pol. Sci. Rev. 76: 738-739.