

THE CRUDE ANALYSIS OF STRATEGIC CHOICES¹

By DANIEL ELLSBERG
The RAND Corporation

Many of the significant implications for U.S. military objectives of specific military choices—such as the introduction of a new weapons systems, a change in basing or deployment, new operational procedures or protective measures—depend upon their impact on the limited set of U.S. and Soviet “decision elements” shown in the following format:

U.S.	Soviet Union	
	Wait	Strike (p)
Wait.....	u_{11}, v_{11}	u_{12}, v_{21}
Strike (q).....	u_{21}, v_{12}	—

We shall not try to define a “game” corresponding to this schema, though the format might suggest that interpretation; rather, it depicts some major, interrelated elements in two concurrent U.S. and Soviet decision problems: whether or not to launch an all-out nuclear attack upon the opponent. “Strike” denotes such an attack; the “Wait” strategy may be interpreted either as a “representative” or as a “best” alternative to Strike. The inclusion of a U.S. Strike strategy does not imply active consideration of such an alternative at any given moment; as a possibility, it is relevant particularly to Soviet calculations, for reasons indicated below.

The u 's and v 's are, respectively, U.S. and Soviet “von Neumann-Morgenstern utilities”² for certain (highly aggregated) consequences of their actions; p and q are, respectively, the U.S. and Soviet subjective probabilities (expectations, estimates of likelihood) that a choice of Wait will encounter an opponent's choice of Strike during a certain time period. Though all of these are subjective variables, they clearly depend upon estimates of objective outcomes under specified circumstances, based upon some form of explicit or tacit “systems analysis.”

¹ This paper is considerably condensed from P-2183, “The Crude Analysis of Strategic Choices” (RAND Corp., Santa Monica, Calif.); P-2183 develops the present approach further and applies it to a number of proposals and propositions in current debate.

² I.e., they indicate not merely order of preference among these outcomes but the decision-maker's preferences among “gambles”: strategies which offer a set of possible outcomes with given subjective probabilities. It is assumed here that utility numbers can be assigned to outcomes so that the decision-maker's actual choices among strategies can be described as maximizing the “mathematical expectation of utility,” the average of these utilities weighted by their respective subjective probabilities.

For purposes of this discussion, the v 's may be regarded as U.S. estimates of Soviet utilities—estimates which are not held with perfect confidence.

Symbols u_{21} and v_{21} are thus, respectively, U.S. and Soviet utility pay-offs for "strike first" outcomes: the consequences of a surprise nuclear attack upon the opponent. Symbols u_{12} and v_{12} are "strike second" pay-offs, reflecting the consequences of being struck first by the opponent. Symbols u_{11} and v_{11} are "no all-out war" pay-offs, corresponding to situations in which neither opponent chooses Strike.³

The precise effects of a change in military "posture," hardware, policy, or plans upon these eight variables (including p and q) are, of course, hard to determine, uncertain, and subject to controversy; nevertheless, rough estimates are often made, and these are, in fact, the basis for most policy recommendations as to choices among military alternatives. Indeed, many such recommendations reflect estimates of effects only upon some subset (e.g., one) of these eight factors. The above schema has the advantage of directing attention at least to these eight, gross consequences of a military change. A typical, major military innovation will affect all of these variables, and in what may be opposing directions for a given or for different military objectives. Within a given set of strategic policy alternatives, "conflicts" may be inescapable; an improvement in terms of one dimension of choice (one military subgoal) may be unavoidably associated with losses with respect to another. Analyses which ignore several of these dimensions are thus likely to be inadequate. Isolated suboptimizing processes which overlook conflicts and "spill-over effects" among related subgoals may end by lowering over-all military security rather than raising it.

While the above highly simplified and abstract schema can by no means be regarded as an adequate model for the comparison of U.S. military alternatives, it may represent a minimum framework which is an advance over that implicit in much current discussion. Assuming that it is possible to estimate the effects of a military innovation (e.g., an airborne alert, the introduction of IRBM's in Europe, a fallout shelter program) upon the factors in this schema,⁴ the question arises: what effects, or complexes of effects, are "good"? For practical purposes the over-all goal of enhancing military security—reducing the likelihood of major losses from the threat or use of enemy military force—must be broken down into military subgoals, a list of specific strategic objectives. Some

³The utility subscripts have been chosen to show corresponding elements in the concurrent, related but separate U.S. and Soviet decision problems; if a "game" formulation were being followed, the subscripts for the v 's would be transposed.

⁴However these estimates are derived and whether or not they are "reliable," the schema can be helpful in deriving their policy implications, in order to test the "consistency" of given policy recommendations with corresponding estimates and objectives.

of these correspond directly to elements in our schema. Thus it is a major U.S. objective to lower p : roughly, to "improve the reliability of deterrence." Likewise, there is the goal of raising u_{12} : improving the strike-second outcome if deterrence should fail. Possible conflicts between these two subgoals are well known. However, by guaranteeing retaliation (lowering v_{21}) it may be possible to lower p greatly, more than compensating for the lower u_{12} which is associated with the low v_{21} . But what is the effect upon p of improving u_{12} , by planning counterforce tactics, or introducing civil defense? To answer this sort of question we must look at the impact not only upon u_{12} but upon all the elements in this framework, for p depends upon the whole configuration of factors in rather a complex way.

To the extent that a Soviet Strike represents a deliberate decision, it must reflect the fact that in Soviet calculations of pay-offs and likelihoods at some moment Strike appeared preferable to its best alternative. The goal of the U.S. "deterrence" policy is to ensure that this never arises: that at all times Strike appears inferior in Soviet calculations to some alternative ("Wait"). In our schema this condition appears equivalently as:

- (1) $V(\text{Wait}) > V(\text{Strike})$, where V is the Soviet utility function; or
- (2) $(1 - q)v_{11} + q \cdot v_{12} - v_{21} > 0$; or
- (3) $(v_{11} - v_{21}) - q(v_{11} - v_{12}) > 0$.

Even though U.S. estimates may indicate that this condition holds at a given moment, p may not be 0; some U.S. uncertainty ($p > 0$) may remain, reflecting: (a) the possibility that U.S. estimates are critically mistaken; (b) the possibility that factors affecting Soviet calculations may change critically within the relevant period; (c) the possibility that Soviet behavior may be non-calculated, impulsive, or erratic, imperfectly co-ordinated, or subject to "unauthorized actions" by subordinates.

Each of these likelihoods is likely to be smaller, the larger the interval, $V(\text{Wait}) - V(\text{Strike})$.⁵ Other things being equal, the "worse" Strike appears relative to its best alternative, then the more likely that the Soviets are "deterred," the more likely that they will stay deterred as pay-offs undergo exogenous shifts, and the more care that Soviet decision-makers will take to avoid accidents, false alarms, hasty decisions, unauthorized actions, or unco-ordinated, unmonitored policies. The size of this interval, then, provides a subcriterion among military choices on the path towards lower p . It is, in effect, an index of the sensitivity of the Soviet decision to "counter-deterrent" shifts in pay-offs (if q is given) such as: (a) a drop in the "no all-out war" outcome v_{11} (due to Soviet losses or

⁵A unit interval having been established by the arbitrary assignment, say, of values 0 and 100 to two specified outcomes.

expectation of losses in a limited war, shifts in prestige, influence or alliances, cold war failures, domestic setbacks or uprisings, political rivalries with third parties); (b) a drop in the Soviet "strike second" outcome v_{12} (due to increased U.S. force size or ability to exploit weaknesses in Soviet warning systems or defenses, or prospect of U.S. "annihilation tactics" in a U.S. first strike); a rise in the Soviet "strike first" outcome v_{21} (a reduction in U.S. "strike second" or retaliatory capability, due to changes either in U.S. or in Soviet posture, procedures, tactics). The larger the interval, $V(\text{Wait}) - V(\text{Strike})$, the larger (in utility terms) the pay-off disturbances required to make Strike appear preferable to Wait. This might be regarded as one index of the reliability of deterrence.

Another important index of this reliability is the sensitivity of the Soviet decision to shifts in q , the Soviet expectation of a U.S. first strike.

To understand why q is relevant at all to the Soviet choice, let us recall the earlier condition of deterrence:

$$V(\text{Wait}) - V(\text{Strike}) = (v_{11} - v_{21}) - q(v_{11} - v_{12}) > 0.$$

Since typically $v_{11} > v_{12}$, it follows that a necessary condition for deterrence is:

$$(v_{11} - v_{21}) > 0, \text{ or } v_{11} > v_{21}.$$

It cannot be taken for granted that this condition will hold; it does not follow automatically from the existence on both sides of nuclear weapons.⁶ But in any case, this condition is not sufficient. Perhaps the most significant aspect of the current strategic balance is that, under typical conditions of technology and posture:⁷

$$v_{21} > v_{12}.$$

It follows that deterrence can fail [$(v_{11} - v_{21}) - q(v_{11} - v_{12}) < 0$] even though $(v_{11} - v_{21})$ is positive and large: if q , the Soviet expectation of a U.S. Strike, is sufficiently great.

An important question is: How high would q have to be to make Strike appear preferable to the Soviets? A threshold value \bar{q} , below which the Soviets would be deterred and above which they would prefer Strike, is given by:

$$(v_{11} - v_{21}) - \bar{q}(v_{11} - v_{12}) = 0, \text{ or } \bar{q} = \frac{v_{11} - v_{21}}{v_{11} - v_{12}}.$$

⁶As Albert Wohlstetter has pointed out, U.S. retaliatory power could be so vulnerable to a Soviet "no warning" attack as to promise less destruction than the Russians have suffered historically, whereas the "no war" outcome could, under abnormal conditions, appear very bad indeed. ("The Delicate Balance of Terror," *Foreign Affairs*, Jan., 1959, p. 222.)

⁷Many of the implications of this relationship between the "strike first" and "strike second" outcomes are exposed in Wohlstetter's brilliant and authoritative article, *op. cit.* In part, the present approach is an attempt to formalize some of the propositions in Wohlstetter's discussion.

We will refer to \bar{q} , that probability of a U.S. Strike which would, with given Soviet pay-offs, make the Soviets indifferent between Strike and Wait, as the "critical risk" of a U.S. Strike. This threshold expectation, defined as a function of Soviet pay-offs, seems a highly significant property of the pay-off structure. Among the most important consequences of military choices is their impact upon this parameter, which serves as an index of the sensitivity of the Soviet decision to their expectation of being struck.

Extreme vulnerability of the U.S. retaliatory force will imply a low Soviet critical risk. It leads to an extreme advantage of the "strike first" over the "strike second" outcome with $(v_{11} - v_{12})$ much greater than $(v_{11} - v_{21})$. With the resulting low \bar{q} , the Soviets would find Strike preferable if they attached even moderate likelihood to a future U.S. Strike. This is clearly an undesirable state of affairs; a Soviet Strike could appear a rational response even to highly ambiguous indications of a U.S. attack, of the sort generated periodically by any warning system. Under the general objective of improving the "reliability of deterrence" it seems desirable to reduce the sensitivity of the Soviet decision to fluctuations in q ; thus, it becomes a subgoal to increase the critical risk, \bar{q} .

The principal method of achieving high \bar{q} —implying that the Soviets will not prefer Strike to Wait unless they are very sure of a U.S. Strike—is to reduce the vulnerability of the U.S. retaliatory force by measures which do not improve markedly the U.S. "strike first" capability; e.g., the replacing of highly vulnerable weapons by Polaris submarines, airborne alert, hardened or land mobile missiles. As v_{21} is lowered relative to v_{12} , a situation is approached in which the Soviets would prefer Wait even if they were certain that the U.S. would attack ($\bar{q} = 1$, corresponding to $v_{21} = v_{12}$).⁸

A further subgoal, towards improving the reliability of deterrence and lowering \bar{p} , is to lower q , the Soviet expectation of a U.S. Strike. Most military choices operate directly upon pay-offs, U.S. and Soviet, with indirect effects on expectations. Changes in U.S. pay-offs will influence q by affecting the Soviet image of the U.S. rational incentives to Strike. Just as \bar{p} corresponds to the U.S. estimate of the reliability of U.S. deterrence, q is essentially the Soviet estimate of the reliability of Soviet deterrence. A way to lower q is to increase, in Soviet eyes, indices of the reliability of Soviet deterrence which are analogous to

⁸ Conceivably, this result might be nailed home by making v_{21} appear worse than v_{12} ; suppose that the Soviets were led to fear U.S. "annihilation tactics" with a large retaliatory force in case of a Soviet first strike, but were also made aware that the U.S. was preparing for a strictly countermilitary campaign, avoiding cities and aiming at quick cessation, if war should arise under any other circumstances. See Herman Kahn, *On Thermonuclear War* (Princeton, 1960), pp. 162-89.

indices of U.S. deterrence: to increase $U(\text{Wait}) - U(\text{Strike})$ in U.S.

calculations; to increase $\tilde{p} = \frac{u_{11} - u_{21}}{u_{11} - u_{12}}$, the U.S. "critical risk"; to

lower p .⁹ This adds two new criteria of choice (lowering p being already included) to our list of military subgoals.

Having presented some apparatus of analysis at this length, there is little space in which to apply it here. Let us consider one example, by now rather familiar. Suppose that, as is frequently done, lowering the Soviet "strike first" outcome v_{21} were taken as the only significant subgoal under the objective of improving the reliability of deterrence; and suppose it were proposed to achieve this by emplacing "soft," fixed, slow-reaction IRBM's in Europe. Like any increase in our inventory of offensive weapons, this move would tend to decrease v_{21} . But only a little; fixed IRBM's are subject to no-warning attack by large numbers of accurate, high-yield Soviet medium-range missiles and bombers, and their existence would have a small or negligible effect on the expected outcome of a well-planned Soviet Strike. Even so, if other effects were ignored, as they often are, the move could seem desirable on the basis of this one criterion.

However, if we ask the impact of this move upon the other factors in our schema, conflicts with other criteria are likely to emerge. The most marked effects of the innovation would probably be: (a) a sharp decrease in v_{12} , the Soviet "strike second" outcome; (b) a sharp increase in u_{21} , the U.S. "strike first" outcome. Neither of these effects, at first glance, might seem undesirable in themselves, to count as "costs." Yet the drop in v_{12} relative to v_{21} would imply a lower Soviet critical risk \tilde{q} ; it would take less assurance than before of a U.S. Strike to make a Soviet Strike seem preferable. And meanwhile, the actual Soviet expectation q might be increased; for the rise in u_{21} relative to u_{12} (which would change negligibly) would mean lower U.S. critical risk \tilde{p} , and, for given p , a reduced interval $U(\text{Wait}) - U(\text{Strike})$, so that Soviet deterrence would appear less reliable than before.¹⁰ Thus, the several criteria we have considered for the reliability of U.S. de-

⁹ An interdependence between p and q emerges here; it has been ably explored under the heading, "The Reciprocal Fear of Surprise Attack," by Thomas Schelling, in *The Strategy of Conflict* (Harvard, 1960), pp. 207-29. I would suggest, without developing the point here, that this interaction is most significant when \tilde{p} and \tilde{q} , the U.S. and Soviet critical risks, are both low.

¹⁰ Herman Kahn has emphasized that such an improvement in u_{21} may significantly improve u_{11} , by deterring the Soviets from such acts short of all-out Strike as might "provoke" a U.S. first strike if the U.S. first strike outcome were sufficiently high. See Kahn, *op. cit.*, pp. 136-44 and *passim*. The objections, which I share, to such a policy of "Type II Deterrence" are too lengthy to discuss here. At any rate, note that this subgoal, if accepted, would in this case conflict with the various criteria of the deterrence of a Soviet Strike (Kahn's "Type I Deterrence").

terrence would indicate that this move might be associated with higher p than before.¹¹ Furthermore, Soviet recognitions of this effect could lead, via Schelling's "reciprocal fear of surprise attack," to higher q and a further upward pressure on p .

These results are to be contrasted to those of the measures mentioned earlier for reducing the vulnerability of the retaliatory force (raising q by reducing v_{21} relative to v_{12}); moreover, a complex of such measures may be designed to raise u_{12} much more sharply than u_{21} , thus providing "second strike insurance" against the failure of deterrence while at the same time raising \bar{p} , increasing the reliability of Soviet deterrence and lowering Soviet fears of attack.

Other specific arms control, civil defense, and active defense measures may be examined in terms of our schema; their implications for the various criteria will depend upon their differential effects upon all of the factors discussed. The discovery, in a particular case, that the implications in terms of several of the criteria (subgoals) conflict is not a failure of the approach; on the contrary, it is a signal of the need for closer analysis in that case, for the weighing of criteria, or for the invention of new alternatives which avoid or alleviate the conflict.

It is clear that this simple framework cannot capture all the complexities of strategic choices. It is in no sense a machine for providing answers; at most, it is a machine for asking useful questions and for preliminary testing of alleged answers. As such, it can be helpful; simple as it is, it is far more flexible and complex than single-variable models implicit in much "literary" discussion. Unfortunately, there has been a historical tendency on the part of policy-makers to reject the aid of abstract frameworks of the present sort on the grounds that they are "too simplistic," and then to make practical decisions on the basis of much cruder, implicit models.

¹¹ This argument follows Wohlstetter, *op. cit.*, pp. 222-30, particularly p. 229.