RANDOMIZATION OF THREATS AND PROMISES

T. C. Schelling

Economics Division

The RAND Corporation

P-1716

June 5, 1959

In the theory of games of pure conflict ("zero-sum" games) randomized strategies play a central role. It may be no exaggeration to say that the potentialities of randomized behavior account for most of the interest in game theory during the past decade and a half.* The essence of randomization in a two-person zero-sum game is to preclude the adversary's gaining intelligence about one's own mode of play -- to prevent his deductive anticipation of how one may make up one's own mind, and to protect oneself from tell-tale regularities of behavior that an adversary might discern or from inadvertent bias in one's choice that an adversary might anticipate. In the games that mix conflict with common interest, however, randomization plays no such principal role, and the role it does play is rather different.

Randomization in the theory of these ("non-zero-sum") games is not mainly concerned with preventing one's strategy from being anticipated. In these games one is often more concerned with making the other player anticipate one's mode of play, and making him anticipate it correctly, than with disguising one's strategy. There may of course be zero-sum components embedded in a larger game. In limited war one may be concerned to communicate rather than to disguise the limits that one proposes to observe, but within those limits may sortie his aircraft in a randomized way to minimize the enemy's tactical intelligence. Again, information samples may be exchanged, or agreements enforced on a sample basis, where neither party can afford to

---

*John von Neumann, speaking of "the fundamental theorem on the existence of good strategies," namely the theorem that all zero-sum games with a finite number of pure strategies have a minimax-maximin equilibrium pair ("solution") if mixed strategies are allowed, said, "...As far as I can see, there could be no theory of games on these bases without that theorem....Throughout the period in question I thought there was nothing worth publishing until the 'minimax theorem' was proved." "Communication on the Borel Notes," Econometrica, Volume XXI, Number 1, January 1953, pp. 124-125.

yield the other full knowledge. Arms-control agreements, for example, might have to be monitored by a sampling technique that yielded each side enough knowledge about the enemy's forces to reveal compliance or non-compliance without yielding so much that the possibility of successful surprise attack on those forces were greatly enhanced.

But the main role of randomization in the traditional literature on non-zero-sum games is a different one. It has been a device to make indivisible objects divisible, or incommensurate objects homogeneous. Their "expected values" are divisible by lottery when the objects themselves are not. We flip coins to see who gets the object, and play "double or nothing" when we cannot make change. We can divide the obligation of citizenship equally by selecting draftees through a lottery, when we want a fraction of the eligibles for a long period of service rather than all of them for a short one.

In this role, randomization is evidently relevant to promises. If the only favors available to be promised are larger than necessary and not divisible, a lottery can scale the expected value of the promise down and reduce the cost to the person making it. An offer to help a person on a large scale in a contingency is somewhat equivalent to offering the certainty of smaller help. (There may be the additional advantage that the contingency is correlated with his need.)

But in this respect a promise is different from a threat. The difference is that a promise is costly when it succeeds, and a threat is costly when it fails. A successful threat is one that is not carried out. If I promise more than I need to as an inducement, and the promise succeeds, I pay more than I needed to. But a threat that is "too big" is likely to be superfluous rather than costly. If I threaten to blow us both to bits when it would have been sufficient to threaten our discomfort, you'll likely still comply; since I have neither to discomfort us nor to kill us, the error costs nothing. If all I had was a grenade to explode in our midst and wished for tear gas instead, I might scale down the grenade to the "size" of a tear-gas bomb by threatening an appropriate percentage chance that the bomb would go off, killing us both, if you failed to comply. But the need to do this is not as clear as in the case of a promise, where any excess in the value promised is so much loss.

The size of the threat can be a problem, though, it costs something to be equipped to make a threat and if bigger threats cost more to make than small ones. If a threat of tear gas in enough, so that I do not need to threaten explosion, and if tear-gas bombs are cheaper than explosive ones, it is better to threaten only the tear gas. But grenades may be cheaper, and then the incentive goes the other way. For many interesting threats the greatest cost is the risk of having to carry it out, and the more ordinary "cost" is not a controlling factor.

THE RISK OF FAILURE

The risk of failure, however, does give an incentive to choose moderate rather than excessive threats. If the only threat that can be made is some horrendous act, one may be tempted to scale it down by attaching it to a lottery device -- by threatening some specified probability that it will be carried out unless compliance is forthcoming, not by committing oneself to the certainty that the jointly painful punishment would be administered.

To illustrate, consider the matrix in Figure 1, in which Column has first choice, followed by Row, but in which Row has the option of making a prior threat to constrain Column's choice. (Interpret X and Y as positive numbers.) On one condition, Row's strategy is clearly to threaten row ii if Column chooses column II. If he makes no threat, Column chooses II knowing that Row will then choose i. Given the threat -- and assuming that Row is committed to it and that Column knows it -- the choice of II yields unattractive outcomes for both of them, and Column can be expected to choose I.

		Column	
		I	II
Row	i	0 1	1 0
	ii	0 0	-X -Y

(Note: Lower-left figure in cell is payoff to Row, upper-right is payoff to Column.)

Figure 1.

The condition is that Row be quite sure that nothing will go wrong! Maybe he completely misjudges Column's payoffs; maybe this particular adversary is drawn from a universe in which nearly everyone, but not quite everyone, has preferences as indicated in the matrix, but a few deviants have a radically different preference system and prefer the lower-right cell to the upper-left one. Alternatively, Row may get himself committed to his threat but fail to communicate it convincingly to Column, so that Column mistakenly ignores the threat, condemning them both to the lower right-hand cell. Again, Column may have arranged a prior commitment through his own choice of II, and failed to communicate it accurately to Row in time for Row to take this into account, or have suffered a disability unknown to Row that eliminates the possibility of I; in that case, Row's own commitment will only guarantee the worst outcome for both players. Whatever the reasons for failure, there is perhaps some probability that the threat will fail. If we take it into account we may have a reason for Row to wish that the "punitive" payoffs in the lower right-hand cell were not quite as unattractive as they are.

If Row is confined to "pure" strategies -- if he must specify his threat or commitment without reference to error or chance -- he can do nothing but wish that the numbers in the lower right-hand cell were not so unattractive. But if he can randomize his threat he can in fact "scale it down" to reduce somewhat the high cost of failure. If, for example, he can commit himself not to a choice of row ii in the event that column II is chosen, but to a 50-50 chance between i and ii in that event, he may still hope to frighten Column into a choice of I while reducing the seriousness of the risk of failure.

We can be more specific. Let P stand for the probability that the threat will fail for any reason whatsoever. (For our present purpose this is an "autonomous" probability, independent of Row's strategy.) Let Row now threaten to choose ii with probability equal to π, in the event Column chooses II. In other words, if Column fails to comply there is a probability of π that Row will choose ii to their mutual discomfort, and of $(1 - \pi)$ that he will choose i to their mutual relief. What value of π should Row choose?

First, how large does π have to be to make the threat effective at all, i.e., to make it effective assuming that it does <u>not</u> fail for any of the autonomous reasons involved in P? This is a question of Column's choice when he is confronted with the risk π. If Column chooses I he gets 0. If he chooses II his expectation is a weighted average of 1 and -X, with weights of $(1 - \pi)$ and π respectively. If this average is less than 0, he is motivated to choose I -- subject to the autonomous probability, P, that for one reason or another he will choose II in spite of his apparent motivation toward I.

The condition for an "effective" threat is thus:*

$$0 > (1 - \pi) - \pi X$$
$$\pi > \frac{1}{1 + X}$$

Second, assume that any threat with π above the floor established by the preceding formula will succeed or fail with probabilities (1-P) and P respectively. If the threat succeeds, Row's payoff is +1. If it fails, his expectation is a weighted average of 0 and -X, the weights being $(1-\pi)$ and π respectively. The expected value of the outcome, then, when the threat is large enough to be effective at all, is given by:

$$(1 - P) + P(0 - \pi Y) = 1 - P - P\pi Y$$

This value is evidently higher, the lower is the value of π. Row should therefore arrange the lowest value of π that he can, that meets the first condition. For a threat to be worthwhile at all -- to have an expected value greater than zero, which is what Row can expect from this particular matrix if he makes no threat -- a value of π must be arranged that meets the condition:

$$1 - P - P\pi Y > 0$$
$$\frac{1 - P}{P} \cdot \frac{1}{Y} > \pi$$

Thus the effective range for π in this example is given by:

$$\frac{1 - P}{P} \cdot \frac{1}{Y} > \pi > \frac{1}{1 + X}$$

And there is no threat at all worth making if there is no room between these two limits, if:

$$\frac{1 - P}{PY} < \frac{1}{1 + X}, \text{ i.e., } \frac{P}{1 - P} > \frac{X + 1}{Y}$$

Only a "fractional" threat -- a threat with π less than 1 -- is wor making if:

*Since the analysis depends only on comparisons of the differences between absolute valuations of the payoffs for the two players separately, no violence is done by adopting, for each player, a scale of measurement that sets his preferred payoff equal to +1 and his next preferred payoff to zero. The full interpretation, then, of the expression $1/(1+X)$, is the ratio of the difference between upper-right and upper-left payoffs (to Column) to the sum of the differences between upper-right and upper-left and between lower-right and upper-left. The simplicity of the formulae thus reflects advantage already taken of this scaling convenience. It takes only one parameter to characterize the relevant relations among three valuations. (In a later problem that involves the lower-left cell, all four payoffs are relevant and a second parameter would be required. That case, however, can be further simplified if the lower-left payoff can be taken equal to one of the others and still illustrate the point; we get less complete knowledge but more zeros and 1's that way.)

$$\frac{1 - P}{PY} < 1$$

$$\frac{P}{1 - P} > \frac{1}{Y}$$

Here is a case, then, in which the fractional threat is superior to the certainty threat, and in which the latter may not be worth making at all when the former may be. The argument hinges on the risk of failure, a risk that has been assumed independent of the size of π itself. This is a somewhat special assumption. If we interpret P as the probability that we have misjudged our adversary and exaggerate his preference for avoiding the lower-right cell, our assumption implies a bimodal distribution of payoffs in the population. It implies that we have either a man whose payoffs are adequately represented by the numbers in our matrix, or a man whose payoffs are so different that no relevant threat -- within the range of values up to $\pi = 1$ -- will dissuade him. If instead we supposed that the ratio of column payoffs in the upper and lower-right-hand cells showed a bell-shaped frequency distribution within the population, and that our particular adversary had been drawn at random, the probability that our threat would succeed would vary directly with the value of π. The probability that a burglar drawn at random from the universe of burglars will be deterred by some specified probability of apprehension and conviction presumably varies directly with the latter probability; the simple model analyzed above treats burglars as divisible into two classes -- those, let us say, who steal for money and are certainly deterred in accordance with the numbers of the matrix, and those who steal for fun and are beyond reach of any threat of the magnitude entered in the lower right-hand cell. On the other hand if our probability of failure reflected, say, a breakdown of communication with the adversary, there might be better reason for supposing the probability of failure to be independent of the particular threat being communicated.

It is interesting to notice that attaching a probability of fulfillment to our threat is, in the above model, substantially equivalent to scaling down the size of the threat more directly. To see this, interpret X in the lower right-hand cell as a fine that will be levied on both Row and Column, or a number of lashes with the whip or days of imprisonment that both will suffer if the threat is fulfilled. If X is the maximum number of dollars, lashes or days that Row can threaten, let π be interpreted as Row's speci ation of what fraction of the maximum permissible penalty is to be exacted; if π is set at .5, for example, both Row and Column receive exactly half their maximum punishments. If we interpret the matrix in this way, and ask what value of π provides the optimum threat from Row's point of view, we go through the same analysis and we reach the same conclusion as before. π is to be as small as possible subject to a minimum value equal to $\frac{1}{1 + X}$. Thus we can interpret π either as a probability of threat fulfillment or as the scale on which the threat is to be certainly carried out. Since the two formulations come to the same thing, and we can interpret π either way, it seems fair to say that _in this case_ the role of randomization is that of making divisible an otherwise too large and indivisible threat, of making possible a "smaller" threat than was otherwise available. (It should be noted though, that to reduce a threat by reducing the probability of its fulfillment reduces the expected value of the outcome proportionately for both players, while a direct reduction in size might not be restricted to proportionate changes for the two parties.)*

THE RISK OF INADVERTENT FULFILLMENT

There is another "cost" element that can motivate a reduced threat. This is the risk that one will fulfill the threat inadvertently, eve i if the

*Randomization may also be integrally related to the arrangement of the threat itself, or be involved in the decision process whether the threatener wishes it or not. So the interpretation of randomization as just a means of manipulating the size of the threat is applicable only in some cases.

adversary does comply with it (or would have complied if the threat hadn't gone off accidentally before he had a chance). The gun that threatens a burglar or hold-up victim may go off accidentally before he has a chance to comply. The dog that threatens to bite trespassers may bite some who do not trespass.

If a hitchiker pulls a gun on the driver of a car and the driver threatens to kill them both unless the hitchiker throws his gun out the window, making his threat by pressing the accelerator to the floor and creating a manifest risk of fatal accident, there is some chance that the accident will occur before the hitchiker has a chance to comprehend the threat and comply. In this case, the risk of accidental fulfillment is an integral part of the threat. The only way one can make the threat is to start fulfilling it. Until the driver speeds up the hitchiker has no reason to believe him; once he does speed up, there is some minimum length of time it takes the hitchiker to comply and the driver to relax his speed. There is therefore an interval, however short it may be, that the risk is present; the risk entailed by the high speed must therefore be one that is small enough to be tolerable to the driver during this initial interval. If instead the car were definitely safe at all speeds under sixty but would certainly skid off the road at exactly sixty and there were no gradations between that carried a moderate risk of accident, the driver could have no incentive to incur a dangerous speed and the hitchiker would know it and not respond to a verbal threat of high speed. It is the possibility of a "fractional threat," a threat that carries the risk but not the certainty of death, that gives the driver anything to work with; but to put it into effect he has to suffer it for some finite period.

If in situations of this kind we suppose -- as is roughly true in the hitchiker case -- that the risk of inadvertent fulfillment is proportionate to the probability, π, that one will fulfill the threat if the adversary does not comply -- if the watchdog's propensity to bite innocent passersby is proportionate to his proclivity to bite those who enter the premises -- a formula is obtained that is not very dissimilar to the one already arrived at. Using the same matrix as before (ignoring this time the probability that a potentially effective threat may fail) and letting $\alpha\pi$ represent the probability of inadvertent fulfillment, the minimum value of π is the same as before. The expected value of the outcome to Row, which must exceed zero if he is to make the threat, is given by the left-hand side of the formula:

$$(1 - \alpha\pi) - \alpha\pi Y > 0$$

$$\frac{1}{\alpha(1 + Y)} > \pi > \frac{1}{1 + X}$$

The optimal threat is again one that barely exceeds the lower limit; there is an upper limit to π that may be less than 1; and, depending on the relative values of X and Y and the "cost" parameter α, it may or may not be possible to find a profitable value for π at all.

RANDOMIZED COMMITMENTS

Having found a rationale for a "fractional threat," we can inquire whether the tactic of "unconditional commitment," too, is one that in certain cases can advantageously be made less than certain. As indicated earlier* a pure commitment -- i.e., a definite commitment to a pure strategy -- is equivalent to "first move" in a two-person, two-move game in which one would otherwise have to move second; it is a means of obtaining the equivalent of first move. We have to relax that interpretation if we suppose that Row,

---

*T. C. Schelling, "An Essay on Bargaining," American Economic Review, Vol. XLVI, No. 3, June 1956, pp. 302 ff.

who has second move in the game but who has the option to commit himself ahead of time, commits himself to a 50-50 chance of choosing row i or ii. To do this one must retain the right to move second, exploiting only the right to commit oneself ahead of time; if one had actually to move first, by a definite choice, the possibility of a randomized commitment would be lost. (The randomized commitment is equivalent to a "first move" determined by a random device with odds set by the player, with the odds but not the actual move known to the other player before his own move.)

The same payoff matrix (Figure 1) can be used to illustrate this situation if we change the rules of the game to permit Row an <u>unconditional</u> commitment prior to Column's choice but not permitting him to make his choice depend on Column's. A firm commitment to ii induces a choice of column I but is wasted because the lower-left cell -- to which Row is now committed -- contains no reward. Row's problem is that he needs row ii to induce Column into I, but he needs row i to profit from I. A compromise can be achieved by a randomized commitment, a commitment to a randomized choice. If Row is committed to flip a coin (50-50 chance) to select i or ii after Column has chosen, Column will choose I as long as X is greater than 1.* In that case Row gets an expected value of .5. If Row sets π (the probability of his choosing ii) at just above $1/(1+X)$ he gets the largest expected value consistent with Column's choice of I. (If Column's payoff in the lower-left cell differs from zero, say .5 or -.5, the formula for optimum value of π differs somewhat.) If Row's payoff in the lower-left cell were -1, no commitment with a greater than 50 per cent chance of ii would serve. And if that payoff were -X or

*I.e., as long as the payoff to Column in the lower-right cell falls short of his payoff in the upper-left as much as the payoff in the upper-right exceeds the upper-left. See the earlier footnote on the scaling of payoffs.

worse, no probability mixture of i and ii would work; any mixture with π large enough to induce column I would be too large to yield Row a positive expected value.

There is another rationale for a fractional commitment. In the case just discussed, it was Row's own preference for the upper cell in I that led him to minimize the value of π. In Figure 2 it is <u>Column's</u> motivation that demands some chance of row i, i.e., a fractional value of π. In this case, a firm commitment to row ii induces Column to choose II; a firm commitment to i induces Column to choose I; no commitment at all leaves Column preferring II; a threat to choose i unless Column chooses I will be ineffective unless Row promises to abstain from choosing ii.

		Column I	Column II
Row	i	2, 4	1, 1
	ii	3, 0	2, 2

Figure 2

In all of these "pure-strategy" cases, Row ends up with a score of 2. He can, however, do slightly better with a mixed commitment. He can, because he and Column are both attracted to column I, disagreeing only over the choice of Row in that column. If he offers Column a 50-50 chance between rows i and ii, Column gets an expected value of 2 in the first column, of 1.5 in the second, and chooses the first. This leaves Row an expected value of 2.5. Since Row has a preference for ii, he wants the highest probability of that row consistent with the need to provide Column with a preference for column I. That is, he wants the largest value of π for which (in the matrix shown):

$$4(1 - \pi) > (1 - \pi) + 2\pi$$

$$3/5 > \pi$$

This particular mixed commitment can be called a combination of a fractional threat with a fractional promise. Row, in effect, "threatens" a relatively high probability of i in the event that II is chosen and "promises" it if I is chosen.

He could do even better if he could make π conditional on Column's choice. Any probability up to .75 for row ii, conditional on a choice of column I, is a sufficient inducement if it is certain that Row will retaliate for column II with row i. But if he is limited to making his threat no worse than his promise is good -- if he has to attach the same probability to both of them -- the upper limit to an effective value of π is .6, with an expected value to Row of 2.6 (and of 1.6 for Column). With a separate π for the promise, the upper limit is .75 for an expected payoff of 2.75 (and only 1.0 for Column).