

THE THEORY AND PRACTICE OF BLACKMAIL¹

Daniel Ellsberg

PREFATORY NOTE, 1968

In March, 1959, while a member of the Society of Fellows, Harvard University, I delivered a series of public lectures in Boston, under the auspices of the Lowell Institute; these lectures were subsequently broadcast over WGBH, an educational FM station. The overall subject of the series was, "The Art of Coercion: A Study of Threats in Economic Conflict and War." The following essay was the opening lecture.

I have reprinted this lecture as it was delivered on March 10, 1959, almost ten years ago. It was before I came to RAND, before I worked on general war "bargaining," or worked on cold wars in the Pentagon, or took part in a hot war in Vietnam: much painful, but mostly relevant, experience. And a lot of relevant history has happened to the world since then—including successful United States coercion in the Cuban missile crisis and disastrous failure against North Vietnam. For these and still other reasons (I have thought more about the problem) I would write fairly differently on this subject as of today: and perhaps I will.

The reason for exhuming this paper as a P at this time is that it received fairly wide distribution in this form eight to ten years ago, resulting in a number of citations in the literature of bargaining theory, resulting in a backlog of requests for reprints. I am happy to fill them, but with the above *caveat*: and the additional one that the simple concepts presented here inevitably look much less novel now than they did in 1959, in particular, after the appearance of applications by Glenn Snyder and others, and much independent work by Tom Schelling. Schelling's "An Essay on Bargaining" [reprinted in this collection] was so stimulating to my work here, and his later work has become so widely familiar, that it seems worth explaining that at the time this was written Schelling had not presented any *formalization* of his pathbreaking notions on bargaining and threats, and it was a major aim of my lectures to accomplish this.

A last caution. At least one reader inferred, and published the inference, that I am not merely an analyst but a fan of blackmail, for individuals and governments. He was mistaken.

PREFATORY NOTE, 1974

I scarcely foresaw in 1959 that in the next fifteen years the government of my own country would be the major practitioner of violent blackmail, or

¹ Any views expressed in this paper are those of the author. They should not be interpreted as reflecting the views of the RAND Corporation or the official opinion or policy of any of its governmental or private research sponsors.

terrorism, in the world. I did not foresee that we would drop 7½ million tons of bombs—four times the total tonnage of World War II—on Indochina between 1965 and 1973, as threat after U.S. threat failed to induce Vietnamese “compliance” with U.S.-picked regimes or with U.S. policies opposing self-determination and social change. Illegitimate coercive violence by states—particularly the United States—continues as in Hitler’s day to provide the bulk of illustrative material for the analyst of blackmail, though officials and media prefer to focus public attention on such relatively small-scale phenomena as bank robberies, hijackings, and political kidnapping by resistance groups, while ignoring the terrorism of states.

Needless to say, every aspect of my own thinking on these matters, conceptual and moral, has come up for reexamination during this period and has evolved greatly since the writing of this essay midway in the Cold War.

Your generals talk of maintaining your position in Berlin with force. That is bluff. If you send in tanks, they will burn and make no mistake about it. If you want war, you can have it, but remember, it will be your war. Our rockets will fly automatically.

NIKITA KHRUSHCHEV, June 23, 1959

Q: What do you think of talk such as this?

THE PRESIDENT: Well, I don’t think anything about it at all. I don’t believe that responsible people should indulge in anything that can be even remotely considered ultimatums or threats. That is not the way to reach peaceful solutions.

PRESIDENT EISENHOWER, July 8, 1959

Like it or not, the language of threat and ultimatum is today—this month, this year—the language of diplomacy. It is, for many, an unfamiliar discourse. For them these lectures, whose subject is the logic and the rhetoric of threats, may provide some guide to current headlines.

A tone of some perplexity seems clear in the President’s closing comments in a press conference in the spring of 1959, in which he had faced questions on Berlin that he had heard earlier over Lebanon, Iraq, and Quemoy: “We are living in a sort of a half world in so many things. We are not—we are not fighting a war, we are not killing each other, we are not going to the ultimate horror. On the other side of the picture, we are not living the kind of normal, what we’d like to call a normal life. . . . We just must not ever be indifferent to what is happening in the world today.”²

And what *is* happening? It wants defining; how shall we name it? Call it blackmail; call it deterrence; call both—as we shall, in these lectures—coercion: the art of influencing the behavior of others by threats. Like those of any art, its techniques, goals, requirements impose certain patterns on the behavior of those pursuing it. Our objective will be to understand some aspects of these patterns in the behavior of people who are engaged in choosing threats, in supporting them, in making them credible and effective.

² *New York Times*, March 5, 1959.

In later lectures, we shall be interested in the operation of a political system in which some are making threats and others are ignoring them, resisting, obeying, or making counterthreats. The curious "half world" to which the President refers—in which we live—is such a system.

To understand the workings of this system is to understand, in some degree, many paradoxical aspects of present-day diplomacy. This includes the most perplexing, awful fact of all: though war has become an "ultimate horror," it has not become impossible.

[Nuclear weapons have one, pre-eminent use in politics: to support threats.] These threats recommend themselves, almost inescapably, as tools of policy not only to expansionist powers but also to status-quo nations seeking only to deter aggression and to preserve an orderly world society. It is perhaps unfortunate that the President's statement, cited at the head of this lecture, is not quite accurate; threats and ultimatums *can* lead to peaceful "solutions" (not, typically, mutually desirable ones), which is why they are used. Yet the gist of the President's remark is valid, and ominous. A mutual process of threat and counterthreat, commitment and countercommitment, can interlock to make all choices dangerous; and to increase the possibility of a war desired by none.

The ability to coerce is a form of power: perhaps the most important form underlying calculations of the "balance of power." Other modes of power important in the relationship of state to citizen—the abilities to educate, to inform, to propagandize, to inculcate habits of obedience or loyalty or respect for authority—are less significant in the relations of sovereign states. The power to *compel* by sheer, overmastering force ("naked power") remains; but its use is limited, practically, to war. [It is the *threat* to compel or to punish—in short, coercion—which is the peacetime tool of diplomacy.] And the ability effectively to threaten is by no means equivalent to the (related) ability actually to compel physically or to punish. In the remainder of this lecture, we shall examine abstractly conditions that bear upon the varying influence of threats.

Abstract patterns of threat-behavior can be identified in many concrete settings: economic bargaining, extortion, the use of lawsuits, the deterrence of crime, normal diplomacy, and conventional war. I shall be drawing most of my examples and applications from the field of national, military blackmail. The analysis I shall propose, nevertheless, derives initially from an interest in economic bargaining, and has its roots in the economic literature on that subject. Economists will recognize the strong influence throughout of game theory,³ of Zeuthen's essay on economic warfare,⁴ and of Thomas Schelling's recent brilliant article, "An Essay on Bargaining."⁵ (My approach, however,

³ The best overall reference on game theory is Duncan Luce and Howard Raiffa, *Games and Decisions* (New York, 1957), partially reprinted herein.

⁴ Frederick Zeuthen, *Problems of Monopoly and Economic Warfare* (London, 1930).

⁵ Thomas C. Schelling, "An Essay on Bargaining," reprinted in this collection.

differs in some fundamental ways from these.) [In this lecture, we shall examine a situation in which a single threatener confronts a relatively passive subject.] The questions we shall ask are very basic ones: what are the data that are relevant to this "blackmailer's" decisions; how are they measured and represented (this last is critical because these variables are subjective); how, fundamentally, do they hang together, how do they interact in their influence on decision.

Let us suppose that I am the coercer, or threatener, or as we shall say, "blackmailer"; you are my subject, or "victim." [The goal of coercion is the goal of all power relationships; my problem is to influence your behavior in accordance with my desires: to make you do what I prefer that you should do.] But unlike most modes of power, coercion is designed to influence not your behavior in general, but what the economist calls your "rational decisions": your choices among alternative actions, insofar as these choices are determined by your own subjective expectations and preferences.

The word "rational" here has a technical and fairly limited meaning. It does not mean "reasonable," in any overall sense. It means only that the choice is in some sense controlled by the actor's current expectations of the outcomes of his actions and by his preferences. Those expectations and preferences themselves may seem highly "unreasonable," or as crazy as they could be. A good deal of "insane" behavior, in other words, might be "rational" in this technical sense, whereas habitual or unreflective responses, though perfectly normal and reasonable, might be classed as "nonrational" behavior. None but "rational" behavior is subject to coercion. Only the man rational in this sense can be blackmailed or deterred.

For coercion operates upon expectations: *your* expectations of *my* behavior. For this to be effective, your expectations must influence your behavior; but that is not enough. First of all, my choice must "make a difference" to you; my actions must affect the outcomes to your actions.

Suppose that I have to give—but also to withhold—something that you want: a job, services, goods, honor, authority. Or—and this turns out to be an equivalent condition—suppose that I can inflict some loss on you, relative to the status quo: by violence, by military action, by using my influence over others who have control over you.

Finally, suppose that I have some means of communicating with you and that I can *change* your expectations, to some extent, of my behavior. Given all these conditions, I can set out to *coerce* you: to influence you to choose the action I prefer you to take, by increasing your expectation that if you do not, I will choose some response leading to an outcome still worse for you than compliance.

Does this sound complicated? If we were to continue purely verbally, it would get more so. We can simplify the argument a great deal by referring to a diagram known as a payoff matrix, which shows the relationships of the variables we are considering: the strategies, and the evaluation of their outcomes.

	Comply	Resist
Accept	90	100
Punish		0

The game of "blackmail" is defined by the following rules. You, the "victim," will pick one of two possible strategies, represented here by the two columns of the matrix, labeled, "Comply" or "Resist." I, the "blackmailer," will pick a row *after* you have chosen; my two possible choices are labeled, "Accept" and "Punish." When we have both chosen strategies, our "payoffs" are shown as the numbers in the matrix at the intersection of that row and that column. *This diagram shows only your payoffs.* The payoff numbers indicate not the concrete, "objective" outcomes of the game—which may be various amounts of money, or exchanges of goods, or levels of violence—but your evaluation of these outcomes, your preferences among them. The numbers are intended to reflect not only your order of preference for the outcomes, but—in a sense we shall consider later—your degree, or intensity of preference among them.

Let us assume that you will choose "rationally." Clearly, in this game it "makes a difference" to you which action I will choose. If you were to choose the strategy, "Comply," you would be guaranteed an outcome to which you have assigned the payoff number, 90. If you pick the strategy, "Resist," you might get 100, or you might get 0, depending on which strategy I chose. (Incidentally, instead of using the phrase, "you might choose the strategy 'Comply' or 'Resist,'" let us just say from now on, "You might Comply," and so on.) Therefore, your expectations of my choice will influence your own choice. If you were *certain* that I would Accept, you would Resist; if you were certain that I would Punish you for Resisting, you would, with these payoffs, Comply.

But what if you are not certain what I will do? I suggest: if the risk of getting 0 is "small enough," you will choose to Resist; if it is "too large" you will prefer the certainty of 90 that you associate with Comply.

Suppose that, before we begin any interaction, or let us say, at a particular stage in our negotiations, I expect you to pick your second strategy, Resist. At that point, in other words, I would expect you to assign a very low likelihood to my choosing Punish: perhaps 0 likelihood. And finally, we will suppose that, for some reason, I would prefer that you choose your first strategy, Comply. Why? Let us defer that question for later.

My problem as a blackmailer is to convince you that I am "too likely" to respond with my second strategy, Punish, for you to accept the *risk* that *your* own second strategy, Resist, would entail. In other words, I must make you believe that 0 is so likely an outcome of Resist that Comply with its "safe" outcome of 90, looks better.

Concretely, these strategies might refer to one of many situations. In a bargaining situation, Resist might correspond to your insistence on your own "last offer," which, if I should Accept, will have the value 100 for you. I threaten that if you do not, instead, Comply with *my* "last offer" to you, which has the value 90 for you, there will be "no deal"; I will break off negotiations,

leaving you with an outcome of 0. In labor negotiations, the threatened punishment for failing to Comply might be a strike or lockout. In the rackets, demolition or physical violence. Between firms, a price war, a lawsuit. In negotiations between states, the issue might be control over a piece of territory: the strategy Punish might cover sanctions from breaking off negotiations, through economic pressure and propaganda, to various levels of war. Obviously, the appropriate payoffs—if they can be estimated meaningfully at all—will be obscure in various degrees for these cases. To give our intuition something to work on in this example, however, let us start by assuming that these particular numbers correspond to *money* outcomes, dollar payoffs.

To consider the strategies abstractly once more, if I can persuade you that I am going to Punish you, my problem is solved; but this may be impossible. In any case, let us say, I cannot rely on doing it. Fortunately for me, it is not necessary. To be effective, my threatened punishment need not be certain, only “sufficiently likely.”

How likely must it be? That is a question that I propose as a contribution to the discussion of bargaining and coercion. For it to be meaningful we must have an operational notion of “degrees of likelihood” in the mind of a player. Game theory, rejecting earlier notions of an “objective probability” assigned to choices by a rational opponent, has avoided the concept of “expectation” or “likelihood” altogether. But this ignores the fact that, typically, a player could tell us quite readily that he believes his opponent to be “more likely” to choose one particular strategy than some other. And this does not mean that he believes his opponent’s choice to be “really” random: For example, chosen by flipping a coin. His statements about “likelihood” reflect merely his own uncertainty, his own “degrees of belief” in his predictions of his opponent’s choice.

If we wish to measure these “likelihoods” on the basis of his observable behavior, not merely his statements, we can offer him various side bets on the choices his opponent will make. By the odds that he insists on before he will bet on various strategies—the odds that make him indifferent between betting on one strategy or another—we can infer, perhaps, that he regards one particular action by his opponent as a “long shot,” another as “almost a sure thing” (that is, he will bet on the first only at very favorable odds; on the second perhaps he will “give” odds). Some experiments have suggested that a player’s choices among these “gambles” may be so consistent that fairly precise estimates of the likelihoods he assigns to various events are possible. In any case, this sort of betting behavior will often give reliable *rough* indications of his different degrees of belief (and I would be willing to make use of the rough data we could get just by *asking* him).

For convenience, I will assume from now on that a player’s state of uncertainty, his expectations, can be represented by a distribution of “subjective probabilities”: numbers obeying the axioms of probabilities but which represent the player’s own degrees of belief. Depending on what operation we choose to measure these, to say that he assigns a probability of .5, or 50

percent, to a particular event (such as an opponent's strategy) will mean either: (a) that he says, or feels, that this event is "as likely as not"; or (b) that we *observe* that he is indifferent between betting *on* this event or betting *against* it: that is, between wagering a fixed sum that the event *will* happen or that it will *not* happen. To say that he assigns more than 50 percent will mean that he regards it as "more likely than not"; or, that he would prefer to stake a fixed sum that it *will* happen than that it will not happen.

[In these terms, the question, "How likely must it seem to you that I will carry it out, for my threat to influence you?" becomes: "What subjective probability, at a minimum, must you assign to my choosing Punish for you to prefer Comply to Resist?"]

If we regard these payoffs as representing money, it seems plausible that there is some likelihood of \$0 that would make you roughly *indifferent* between the certainty of \$90 and a gamble that offered you \$0 with that probability and \$100 otherwise. Thus, if you assigned that probability to Punish, you would be indifferent between Resist and Comply. If you thought the risk that I would Punish were *greater* than that, you would Comply; if you assigned *less* than that likelihood to my carrying out my threat, you would Resist. That is, you would accept this (small) risk of getting \$0, in hopes of getting \$100 rather than the \$90 you could get by complying. This hypothetical probability, whatever it is, then forms a threshold: I will refer to it from now on as your critical risk, defined for these particular payoffs. Since it represents the maximum risk of punishment that you will accept, in choosing to Resist, we might term it your "willingness to resist": keeping in mind that it is measured by this hypothetical threshold probability—which in turn, as we shall see, depends on your payoffs.⁶

Now we can imagine you making your choice by comparing your actual expectation, in the end, with your critical risk; you Comply if you decide that the "actual risk" is higher than the critical level; otherwise you Resist my threats, refusing to comply. [And my problem as a blackmailer is to ensure—by actions that either change your payoffs, hence your critical risk, or that increase your expectation of punishment—that your estimate of the actual risk is greater than the critical risk. How to do this is, of course, the heart of the blackmailer's art.] ✕

I do not mean to imply that either the blackmailer or his victim sees the problem in just this way, or that they make any such comparisons consciously. I propose merely to interpret, explain, their behavior in terms of these concepts. And even if I, as a blackmailer, did imagine that you have some threshold probability which I must surpass, neither I, nor any third party—or for that matter, even you—could estimate that threshold at all precisely.

⁶ Frederik Zeuthen has an equivalent threshold concept in his essay, "Economic Warfare," reprinted herein. He applies it to a bilateral bargaining situation, a context we will consider in later lectures. I am indebted to Harvey Wagner for this important reference.

Still, let us see whether our intuition cannot provide some very rough hint as to where your critical risk might lie with the payoffs we have assumed.

Let us suppose, to begin, that the numbers in the matrix represent money payoffs. If the victim complies, he gets \$90. If he resists, he may do \$10 better; he can get \$100 if I fail to carry out my threat. On the other hand, he may do \$90 worse; he will get \$0 if I do carry out my threat. "Resist" thus has the character of an "all-or-none" bet. He will resist if he is certain that I will not carry out my threat; but he will comply if he assigns more than some particular, roughly defined probability to my carrying out the threat. Before he will choose to comply, how sure does he have to be that I will do what I say? How sure would you have to be?

Is your answer like mine: *not very sure*? Then your "critical risk"—as mine would be—is *low* in this situation. If your intuition does not turn up an answer immediately, let us approach one more gradually. Would you decide to comply if you were 90 percent sure that I would punish you for resisting (that is, 90 percent sure that you would get 0)? If so, your critical risk is below 90 percent. Would you still take the sure \$90 if you thought it merely "more likely than not" that I would carry out my threat? Then your critical risk is below 50 percent. (You will not "accept" even a .5 probability of punishment, when your alternative is a sure \$90).

What would you choose if you assigned only a 5 percent chance to my punishing you? If you would then quickly choose Resist, accepting the chance of punishment in return for the hope of gain, we can infer that your threshold probability, your critical risk, is greater than 5 percent, and so on. Before too long, no doubt, your answers would become halting and inconsistent; we would be left with a range of, say, 5- to 20 percent covering your elusive critical risk: that minimum probability of punishment that would shift you from Refuse to Comply. This critical risk represents the minimum "credibility" that a blackmailer must achieve for his threat to influence you. Facing this empirical data, a conservative blackmailer would aim at achieving at least a 20 percent likelihood in your mind that you would punish him for refusing.

[Incidentally, if you happened to be the sort who followed the rule of thumb: choose that action with the highest "expectation" of money, that is, with the highest weighted average of money payoff, then I could compute your critical risk precisely.] In this case it would be 10 percent. If you assigned 10 percent probability to \$0 and 90 percent to \$100, given your second strategy, you would be just indifferent between that strategy, which would then offer an expected payoff of \$90, and your first strategy. If you assigned more than 10 percent, you would comply; if less, refuse. (There is a simple way to compute this result, which will be presented later.)

[But (in all cases) such precise calculation is thoroughly unrealistic; the usefulness of this approach in no way hinges on achieving such spurious accuracy. At most, what we will hope to agree on is a judgment of the form:] "The threshold is *low*; the victim does not need to be *very* sure of punishment before he will comply." Or an even cruder judgment can still be useful: "Low

or not, the critical risk in this case is *lower* than in some other": That is, lower than in the following case:

	Comply	Resist
Accept	10	100
Punish		0

Here, in money terms, only my "offer" has changed: Comply now gives you 10. Refuse still offers, for you, a possible \$100: a possible \$0. Yet, the overall *structure* of the payoff matrix has changed. Has anything happened to your "willingness to resist": your critical risk? By complying now you accept the certainty of \$10. If you refuse, you will sacrifice that \$10, if I carry out my threat; but you could gain \$90, if I should fail to do so. In the earlier case, let us say, you would have complied if you had assigned as much as one-fourth, even one-fifth, likelihood to my punishment. Would you still?

My own intuition gives me an answer. It seems to me that you will "need" to be more sure than before that I will carry out my threat before you will comply; that is, whatever your critical risk was before, I will guess that it is *higher* now. [Now that I offer you less for compliance, I must manage to make my threat more credible than it had to be before.]

In fact, with these payoffs, I would be willing to go further and guess not only that your critical risk now is higher, but also that it is fairly *high*. Not 27 percent, not 93 percent; just, *high*. My estimates, indeed, may never get more precise than that; but that may be just close enough.

[In practice, of course, your choice will not depend only on the *money* outcomes; your willingness to take risks involving these outcomes will reflect many other considerations.] For example, if you defined the messages I was sending you as "blackmail," you might find "submitting to blackmail" so distasteful that you would Refuse unless you were "almost certain" that I would carry out my punishment, or perhaps even if you were certain. In the analysis, your payoffs should reflect such feelings; they imply that the payoff to "compliance" would appear very low, perhaps even lower than that of "punishment." Again, even a small loss, or a small gain, may weigh heavily with you. For this analysis, we want payoff numbers that indicate to us just how you *would* gamble on these outcomes. The outcomes will typically not be money, or money alone; there may be no plausible "objective" index of the "size" of the outcome. But we are not interested in "size" in any objective sense; all we are interested in is the player's subjective, relative evaluation of the outcome: specifically, what bets he will take, what he will reject, in choosing among gambles that offer these outcomes with various probabilities. If we can actually observe his choices among such gambles—fine. As L. J. Savage has shown,⁷ following the work of Frank Ramsey,⁸ it may be possible to estimate both payoffs and probabilities from this pattern of choices. But I

⁷ L. J. Savage, *The Foundations of Statistics* (New York, 1954).

⁸ Frank Ramsey, "Truth and Probability" (1926) in *The Foundations of Mathematics* (London, 1931).

suspect that we may be able to make fairly reliable, though rough, guesses about a player's relative payoffs by inferring his "degree of preference" from many other types of evidence besides his betting behavior. An actual black-mailer, after all, must make such guesses; and he will; and at least in some, sufficiently "obvious" cases, I think we can predict roughly his answers.

What went through the mind of the bank teller in New York last December, as he read the note that a "little old lady" pushed through his window?⁹ "I have acid in a glass," the note said, "and if you don't give me what I want I'll splash it on you." He looked up, saw about ninety customers in the bank, a grey-haired lady in a brown cloth coat before his window, and on his ledge, a six-ounce water glass with a colorless liquid in it. He returned to the note, and read: "I have two men in here. I'll throw the acid in your face and somebody will get shot. Hurry. Put all the fives, tens and twenties in this bag." He complied.¹⁰

Why? A matter of reflex? Try that on the reflexes of your own teller. Because he was quite certain that the glass contained acid, and that she would throw it at him? An unlikely estimate. (It turned out later that she had stolen the glass from a cafeteria and filled it with tapwater.) Did he obey because he hastily calculated that his critical risk was 14 percent, and the actual risk she presented was 23 percent? No, and yet it would not require much reflection on his part to conclude that there was *some* chance that she was telling the truth; and looking into his heart, he might well have concluded that he did not need to be *very* sure.

The little old lady had to face beforehand the problem of estimating these factors. As predictions go in the social sciences, hers were pretty good.

We have more data on the mental processes of the next teller to face the lady. This was in a bank one block north, the following month; the operation was virtually identical, the note in this case reading: "Put \$5000 in \$5, \$10, and \$20 bills in bag. There is acid in the glass. You wouldn't want it in your face. Don't try to warn anybody until I am out of here. Oh yes, I am being guarded by two guns. Let's not let your customers or co-workers get hurt."

The newspaper account continues: "Mr. _____'s first impulse when he read the note was to duck behind the counter and set off the alarm but he reconsidered when he looked up and saw the water glass. 'She tilted it just to show me that there was something in it,' he recalled later. He put his right hand up in front of his face and with his left began cramming bills into the paper bag."¹¹

As it happened, the teller's caution led him to hand the paper bag over the glass partition instead of unlocking the grille; this lapse from normal procedure

⁹ *New York Herald Tribune*, December 4, 1958.

¹⁰ This robbery was accomplished at lunch hour on one of the busiest corners in the world: Seventh Avenue and 34th Street, across the street from Macy's. The lady walked out, with glass in hand, past an armed guard; she had not uttered a word. Her exit was followed closely by the entrance of two armed guards, two uniformed patrolmen with drawn weapons, and two FBI agents who happened to be on the corner.

¹¹ *New York Herald Tribune*, January 3, 1959.

was noticed by a vice-president standing nearby, who had also seen the note being passed. He shouted, "Grab that woman!" and the lady surrendered quietly.¹²

[If we accept the notion that we can make some meaningful estimates of "critical risks" in different situations, we can ask: On what does the critical risk depend? What factors make it go up or down? This is of interest to the blackmailer, since, other things being equal, he will prefer a situation in which his threat does not "need" to be so credible. The general answer is that it is the whole structure of the victim's payoffs that determines the critical risk.] In this simple case, as the possible gains from Resisting, relative to the outcomes to Complying, increase relative to the possible losses, then the critical risk, the "willingness to Resist," goes up. For example, I can reduce your critical risk by threatening a worse punishment than before. Or, as we have already suggested, your critical risk goes up the worse the "offer" I make to you: That is, the worse your outcome to Comply, relative to the payoffs to Resist.

This last point was illustrated in an interesting bargaining session in a Chemical Corn Exchange Bank eight blocks away from the first in which the little old lady scored. The note, typed on the back of a check, read: "Have a grenade in my hand and two more. I will use them. Didn't come here to get caught. Take a look. Toss in your and your neighbor's money. Put \$5000 in bundled 20 and 10s in coin bag. Count it. Don't stuff it. Keep quiet until I'm out the door or this place will look like a battlefield."¹³

To this the man added verbally: "I'm not fooling." The teller told police the man was clutching "something" in his left hand and he "believed it was a grenade." The teller counted out \$5,000 from a cash supply of \$50,000 and put it in the bag. Then occurred a little exchange that suggests that, given the teller's actual expectation that the man was holding a grenade and would use it to blow up the teller (and himself), he regarded the actual risk as "too high" to argue with \$5,000 at stake, but not high enough when the "offer" was very much worse.

Eying the remaining \$45,000 in the pile, the robber asked: "What about the rest?" According to the newspaper account, "Mr. ——— told him: 'That's all you asked for, \$5,000,' and handed over the bag containing \$4,000 in \$20 bills and the rest in \$10 bills." The man departed without further argument, passing the guard and crossing Fifth Avenue in front of the Public Library.

A crucial element in the success of all these robberies was that, confronted by the threat of acid or grenade (combined with modest demands), the tellers'

¹² She was eventually convicted of the two robberies (when confronted by her first victim, she volunteered, "Oh yes, that's the boy I held up last time") without ever having given her address (or, some suspected, her correct name). As for the \$3,420 from the first robbery, she informed police: "A few seconds after leaving the bank, the paper bag with the money was jostled from my hands and fell to the street. A man picked it up and tried to hand it to me but I was so nervous I said, 'It doesn't belong to me' and walked on."

Police described the story as "fantastic."

¹³ *New York Herald Tribune*, November 4, 1958.

critical risks ("willingness to resist") were *low*. Now, any one of us might have predicted, as did the robbers, that this would be the case. This does not necessarily mean that you should give up a good, alternative means of livelihood to make your fortune with a glass of water and a paper bag. There is more to the blackmailer's art than estimating the victim's critical risk. How does he (or she) manage to ensure that his victim will estimate the "actual risk" as *higher* than his critical level? How does he do this with sufficiently high confidence that he is willing to gamble on his own success? What are the obstacles in his way, the limitations on his ability to influence his victim's expectations? These are the central problems in coercion.

They involve, primarily, influencing your expectations. Well, how do you form your expectations of my behavior in the first place? How would you arrive at them if I were doing nothing to influence you?

You can go by past experience, if you know any that seems relevant. How have I behaved in "games" with others? With you? Do I "tend" to carry out my threats, my predictions? Or do I bluff? How do people "like me" typically behave—on the basis of your experience, or what you read in the papers? If I know "the record" you are likely to be looking at, I can guess at your initial expectations. I may be able to fill in the record a bit for you, underline certain parts, lie about it; but on the whole there is little I can do to manipulate your expectations by changing the record.

With or without experience to go on, the victim has another basis for expectations: his estimates of the blackmailer's payoffs. If you regard me as "rational," in our limited sense, then your question, "How is he likely to respond to my actions?" translates into: "What will he *want* to do? What will it pay him to do? Which will he see as the *best* choice for him to make, given my choice?" You may have no idea what my payoffs are, in which case these questions go unanswered. At best, your estimates will be rough and uncertain, perhaps even more than my estimates of your payoffs, since it is typical for the blackmailer to know a good deal more about his victim than the victim knows about him. Still, you usually will make some assumption, perhaps very well informed, about my payoffs; and when you do, it will influence your expectations.

[And this can be the blackmailer's primary obstacle, the crucial limitation on his ability to influence. The effect of his own payoffs on the victim's expectations will often be the factor that he must change, or counteract, or surmount, if he is to succeed.] For typically, if the victim should fail to "obey," it would be *costly* for the blackmailer to carry out his threatened punishment. It would not, in general, give him his best outcome under the circumstances: that is, given the victim's actual, rebellious choice. Which is to say, it would be *irrational* to carry out the threat. It would mean, for the blackmailer, deliberately passing up an outcome that he preferred; he could, if he chose, do better. Why wouldn't he? That question is bound to occur to the victim; and the blackmailer knows it.

We can illustrate this problem with our payoff matrix by adding to the

diagram the blackmailer's payoffs. These numbers will appear to the left of the victim's corresponding payoff; from now on, the first number in the box at the intersection of a row and column signifies the payoff to the "row" player (the blackmailer), the second the payoff to the "column" player (the victim). For example:

	Comply	Resist
Accept	100, 10	50, 100
Punish		0, 0

It is a while before this type of diagram begins to look as simple as it really is; the trick is to learn to look only at the left-hand payoff in each pair—by rows—when we are interested in the blackmailer's payoffs, and at the right-hand payoffs—by columns—when we are discussing the victim's payoffs. [We will assume that there is no basis whatever for comparing the victim's payoffs directly to the blackmailer's payoffs, or vice versa. One might be expressed in dollars, the other in rubles, with no given exchange rate; or one might be dollars, the other "utilities" (measured by some psychological test). A given payoff for the victim, say, has meaning only in relation to the other payoffs *for the victim*; there is no basis for saying that it is "more" or "less" than any given payoff of the blackmailer.¹⁴ Thus, if the blackmailer Accepts and the victim Resists, there is no basis in the numbers above for saying that the victim's payoff is better, in any sense, than that of the blackmailer (to make the point in familiar, though not quite conclusive, terms: the victim's payoff might be expressed in pennies, the blackmailer's in dollars).

In this example, if the victim Resists and the blackmailer carries out his threat to Punish, the victim gets 0, instead of the 10 he could have had with certainty if he had chosen Comply. But the punishment is costly also for the blackmailer; his payoff, from choosing Punish, is 0 instead of the 50 that he could have had (given the victim's prior choice; the blackmailer, remember, always moves second) by choosing Accept. Can he convince the victim that, if the occasion arose, he would pick 0 instead of 50?

[The first point to make—following from all our earlier discussion—is that he need not make the victim *certain* that he would do this. The victim must believe merely that the likelihood of the threatened action is greater than his own "critical risk," a factor which depends only upon the structure of his own payoffs.¹⁵] The question we now face is: how does the blackmailer make his

¹⁴ I am underlining this point so much because this is a major distinction between the approach I am proposing in these lectures and many earlier treatments of threat-situations. For example, Luce and Raiffa, in their discussion of game-theoretical approaches, several times mention that it is possible to judge the effectiveness of threats only if comparisons may be made between the players' payoffs. In other words, they, and other writers, have stressed that it is necessary for the threatener to be able to say, "This would hurt you more than it would me." Many critics have doubted that it is possible to make such comparisons, on a meaningful, empirical basis; and my own approach convinces me that it is not *necessary* to do so. At any rate, we will not.

¹⁵ I have not been drawing comparisons between the approach presented here and alternative, earlier approaches. This is a convenient place to remark that the above point is, so far as I am aware, a novelty of the present analysis.

threat appear even *slightly* likely, against the evidence of his own payoffs? For with the payoffs as given, it would clearly be irrational for the blackmailer to carry out the threatened punishment, if the victim, after all, failed to Comply. And this example is not peculiar. It is not the exception but the rule for a threat to have a certain built-in implausibility, being costly—or irrational—for the threatener to carry out. His efforts to overcome this barrier to belief account for the most characteristic, and paradoxical, features of threat-behavior. (I owe this conception of the central problem of threat-behavior to Thomas Schelling's pathbreaking article, "An Essay on Bargaining," reprinted herein.)

A good deal of our subsequent discussion will be devoted to illustrating and analyzing the techniques whereby a threatener (including a victim making counterthreats) attempts to make it sufficiently plausible that he will carry out a costly threatened action. There seem to be four main approaches, of which the first two, in particular, have been discussed exhaustively in Schelling's article, to which I refer the reader for extended treatment.

First, the blackmailer can voluntarily but irreversibly give up his freedom of choice;¹⁶ he can make it impossible for himself not to carry out his threat. If he can in some way bind his own hands, destroy his alternatives, he may be able to "make it true" that he would carry out his threat, for the reason that he would have no choice. For example, he might irrevocably give control of the punishment strategy to a subordinate, an ally, an agent who would actually be more likely to carry it out (presumably because his payoffs have a different structure). United States relations with Chiang on the Quemoy issue might be pertinent here, or the proposal to give NATO nations independent control of nuclear weapons.

This might be represented in the matrix by deleting certain rows, or entries, symbolizing: (a) that these have been eliminated as feasible alternatives for the blackmailer; and (b) that his opponent recognizes this. Thus, in our example, the blackmailer might be able to "strike out" the first row, perhaps thus "compelling" the victim to believe that the consequence of Resist must be 0.

	Comply	Resist
Accept	100, 10	50, 100
Punish		0, 0

Here we find the blackmailer striving to achieve his goal by eliminating opportunities, contracting his set of alternatives, although discussions of bargaining often rule this out axiomatically. This was the type of possibility, emphasized by Schelling, that gave his article a paradoxical flavor for most readers. The resolution of the paradox is that the blackmailer, by his tactic of "tying one hand behind him,"¹⁷ hopes for a favorable effect on his opponent's expectations.

¹⁶ Schelling, "An Essay on Bargaining," reprinted herein.

¹⁷ At this point we might look again at the wording of Khrushchev's statement to

(2) [The same flavor of paradox, and the same explanation, attaches to the next type of tactic. It may actually be rather rare that a blackmailer can tie his hands irrevocably, quite literally destroy alternatives. But even when he cannot make the actions in question (that is, failure to punish) impossible, he may be able to make them *costly*. Although Schelling does not emphasize the distinction in that article, most of the examples of "commitment" in his essay fall into this category. The player binds himself to incur certain costs or penalties or to forego certain advantages if he should fail to carry out a pledge. Thus he reduces his own payoff incentives to break the pledge—perhaps to the point where it would become irrational to break it—thereby hoping to make his pledged action seem more likely to his opponent.

We can represent this behavior neatly in our formal model by allowing the player the opportunity to *lower his own payoffs*; this "move" serves to formalize most of the behavior examined by Schelling, behavior which is ignored or excluded in most game-theoretical discussion and which tends to appear puzzling or perverse in actual experience.¹⁸

Thus, the blackmailer may seek to *change* his payoffs from those presumed in our example to the following:

	Comply	Resist
Accept	100, 10	-1, 100
Punish		0, 0

Here he has done nothing but *worsen* one of his own payoffs; yet he may have improved considerably his chances for winning a favorable outcome! He has made it plausible that he will carry out his threat by making it rational to do so.¹⁹

Concretely, the blackmailer may make a contractual agreement with a third party to choose Punish if the victim chooses Resist, making himself liable to forfeit or penalty or suit if he fails to carry out this action. Or he can stake his honor, his prestige, his reputation for honesty—if he has any of these—on carrying out this prediction. These new obligations are then symbolized in our model as subtractions from his previously determined payoff to choosing Accept when the victim chooses Resist (that is, the payoff for failure to punish). If he can actually lower that payoff below that for Punish (and even, to some extent, if he cannot go that far) he will have removed (or

Harriman, cited at the head of the lecture: "Our rockets will fly *automatically*." "And his colleagues around the table," Harriman continues, "chorused the word, 'Automatically.'" *Life*, July 13, 1959, p. 33.

¹⁸ In Schelling's most recent publication on this subject, "The Strategy of Conflict: Prospectus for a Reorientation of Game Theory," *Journal of Conflict Resolution*, II, no. 3 (Sept., 1958), 240, he has proposed independently this same formalization of his earlier work.

¹⁹ It might seem natural to infer from this that the victim is subsequently *certain* to Comply, since, it might be argued, he will assume that the blackmailer is now certain to carry out his threat. However, there are good reasons to be discussed for avoiding this assumption. The fact is that this tactic does have risks for the blackmailer; it does not guarantee him a win.

reduced) his evident incentive to back down from the threat, which presumably becomes more convincing to the opponent. Thus the importance of honor, reputation, prestige, to a blackmailer. They can be pawned. They can be wagered, risked, put up as security; they are something to lose, which can make more credible choices designed to preserve them.

Or the stakes can be more tangible. Chiang made a threatened defense of Quemoy against invasion far more credible, though this may not have been his direct intention, by stationing one-third of his troops on the island: thus "making it true" that he had a great deal to lose by accepting the loss of the island.

[These two types of tactics are basic, but they have certain drawbacks in making credible what amount to *suicidal* threats. Whose honor, whose reputation for honesty is so great that to wager it would make it actually *rational* to carry out such a threat? And who, with such issues at stake, would really bind himself irrevocably to carry out a suicidal punishment?]

The problem is not an academic one, for precisely these questions have been raised repeatedly about the use of the deterrent threat of massive retaliation. As Secretary of State Dulles described the policy,²⁰ "The basic decision was to depend primarily upon a great capacity to retaliate, instantly, by means and at places of our choosing." No one denied that we had this capacity; the doubt was that we could make an opponent believe that we would use it, on occasions when it would obviously not be the optimum move to make. And that—in a world in which the Russians also possessed atomic weapons and a long-range air force—seemed to cover just about every occasion short of an all-out attack on the United States.

In 1952 Dulles had stated: "The only way to stop prospective aggressors is to convince them in advance that if they commit aggression, they will be subjected to retaliatory blows so costly that their aggression will not be a profitable operation."²¹ But *how* convince them, when it is evident that retaliation might be no less costly for *us*?

The point is obvious if we attempt to sketch a payoff matrix. First, say, the Soviet payoffs might have this character:

		SU	
		Status quo	Aggression
US	Accept	0	10
	Retaliate		-1000

Their two (column) strategies are labeled Status quo, and (limited) Aggression; the latter refers to some expansionist move which does not include an attack upon America. The United States strategy of Retaliate implies massive retaliation on Russia, using SAC and thermonuclear weapons.

Clearly, if the Soviet Union were certain that we would retaliate thus, they

²⁰ Department of State *Bulletin*, XXX (January 25, 1954), 107-110.

²¹ *New York Herald Tribune*, May 16, 1952. In this formulation, the essential similarity shows up clearly between the "deterrent" threat and the "blackmail" we have been discussing till now.

would refrain from their limited aggression. But would they be certain? It is not enough, the critics of the massive retaliation strategy pointed out, for the United States simply to tell them we would retaliate; for consider the corresponding United States payoffs (which we show to the left of the Soviet Union payoffs):

	SU	
	Status quo	Aggression
Accept	0, 0	-10, 10
Retaliate		-1000, -1000

If the Soviet Union chose their limited aggression and we "accepted," we would be worse off than the status quo; we would have suffered a loss of -10. But if we carried out the threat of a nuclear attack on Russia, the loss would be immeasurably greater: -1000. Would the United States take that choice?

"I think the answer is clearly no," said Dean Acheson recently, after putting this question to himself.²² "To invoke our 'great capacity to retaliate' is to insure an experience by us of the Soviet's 'great capacity to retaliate.'"

We have already seen that the general problem of credibility is by no means peculiar to threats of massive retaliation. On the other hand, the tactics considered so far—binding oneself irrevocably, or putting up forfeits—do not seem too helpful to the threatener in this case. Is it credible that we would be willing really to cross out the first row in that matrix? Is our honor (or anything else) really so valuable that, by pledging it against our threat to retaliate, we would reduce -10 (in the upper right-hand corner) to -1010 (thus making retaliation "rational")? "No," the critics of massive retaliation have generally answered; and there is valuable insight in their reaction.

Still, this is not the whole story. These critics have often weakened their case—from the point of view of this discussion—by appearing to assume that the threatened retaliation must be perfectly certain to be effective at all. In his book *Power and Diplomacy* (written a year before the 1958 Quemoy crisis), Acheson asked:²³ "Do any of us seriously believe that an American government would take the position that an attack on Quemoy would involve the destruction of Peiping or Moscow, or both, and of New York? "The answer is, of course, that the threat is not credible."

And it is "the essence of deterrence," he asserts earlier, "that the threat should be credible."²⁴

These assertions suggest the existence of just two classes of threats: "credible" or "incredible." Now, we can make finer distinctions than that; and we must, for a real understanding of this problem. [The essence of deterrence (as of blackmail) is that the threat should be credible enough.] To put the problem that way is to suggest the question: just how credible must this particular

²² Dean Acheson, "Wishing Won't Hold Berlin," *Saturday Evening Post*, March 7, 1959, p. 85.

²³ Dean Acheson, *Power and Diplomacy* (Cambridge, 1958), p. 51 (italics his).

²⁴ *Ibid.*, p. 47.

threat be? And a glance at plausible payoffs for the massive retaliation threat shows up another peculiarity of that situation; while the threat seems unlikely to be very credible, it would not appear that it *had* to be very credible.

Consider our hypothetical payoffs. Soviet aggression offers them a possible gain of 10 over the status quo (if the United States does not retaliate) but a possible loss of -1000 (if we do). Given the United States payoffs, it would clearly be hard for the United States to make its threat seem, say, "as likely as not." But given the Soviet Union payoffs, it is hard to believe that they would *need* to assign that much likelihood, for them to be deterred. With payoffs like these, in other words, it seems plausible (to me, at any rate) that the Soviet "critical risk" would be well under 50 percent.

This is not an argument in favor of relying on "low-likelihood" threats. Acheson's judgment that the threat of massive retaliation might not be credible *enough* (to rephrase his statement in our terms) in a particular case is undoubtedly sound. But unless we face the full complexities of the problem, we cannot understand the incentives that can lead reasonable men to commit themselves to relatively implausible threats (on the conceivably well-founded hope that they will prove plausible enough). Nor would we prepare ourselves for the possibility that we might have to *face* such threats from our opponents.

Given, then, that our threatener may not have to make his threat very credible (his victim's critical risk may be low), but that the tactics we discussed earlier may not suffice to make it credible at all: there remain two more classes of tactics that may be effective. First, as a blackmailer, I may create and exploit uncertainty in my opponent's mind as to my true payoffs: make him unsure in his predictions of the actions that would truly be rational for me. Finally, I can appear *irrational*; for which purpose, of course, it helps to *be irrational.* Irrational, perhaps, in being erratic, inconsistent, unpredictable; or again, in having abnormal, "unreasonable" expectations or preferences.

I shall discuss these methods later in detail in an examination of Hitler's diplomacy: "The Political Uses of Madness." Meanwhile, we might note that they seem highly relevant to the bank robbery examples mentioned earlier. In those cases, a low critical risk was a common element on the part of the victims; at the same time, an obviously "reasonable" threatener would have had trouble in making the threats in question seem plausible in the slightest. The teller might well reason: "If I were standing out there, I would *never* throw the acid, or drop the grenade, no matter what happened." But his next thought is bound to be: "*I* would never be standing out there with a grenade in the first place." The very presence of these people in the bank, pushing homicidal notes through the window in full presence of customers and guards, is very impressive warning to the teller not to rely on inferring their payoffs and behavior patterns from his own.

This sort of consideration will usually create only a small uncertainty in the victim's mind; but given a low critical risk, that may be enough. And critical risks can be very low. About the same period that the little old lady was operating on Seventh Avenue, a young electrical engineer held up two Chemi-

cal Corn Exchange Banks in succession (one in Penn Station; one on Broadway near 38th) *with a comb*. Each time he asked for \$5,000; the note in the second case, on the back of a deposit slip, advised: "Look at the gun. Don't try to signal. Don't try to press alarm. Put \$5,000 in 100s, 20s, and 10s in this envelope. Keep quiet till I'm out."²⁵ The instruction to look at the gun was a nice touch, since he did not have one. As the newspaper put it: "A comb stuck in _____'s belt, high across his chest, appeared to Mr. _____ to be the handle of a pistol. He immediately counted out \$7,060 and thrust it into the brown manila envelope _____ pushed toward him." (In this case, the robber was caught after a chase; but, it turned out, his *three* earlier comb heists had been successful.)

Notice the recurrence of the \$5,000 figure in the successful "offer"; interesting, too, that the man with the bag of grenades had selected a Chemical Corn Exchange Bank. Had the word gotten out about the critical risks of Chemical Corn Exchange tellers? "ask for \$5,000 and you're in . . . point anything at them—cookies, milk bottles, rolled newspapers. . . ?"²⁶ Their "willingness to resist" seems on a par with that of an American POW described in a recent book on the POW's in Korea. Explaining why he had given information after thirty-five minutes of mild questioning, he said: "They said they had ways of making me talk, so I talked."

Incidentally, notice the rich array of possible explanations our discussion suggests for this man's behavior. His critical risk might have been very low; either because (a) torture had a very great negative payoff for him, or (b) the information requested represented to him a very small "demand." Or he might have regarded the actual risk of punishment as very high, because of clues in the environment, or shared experience, or conservative assumption. The Army's categorization of its POW's is less fine-grained; they decided he was a "coward."

It is a peculiarity of thermonuclear threats that they make cowards of practically everyone. Very few are the objectives that would seem worth—to anyone—a high probability of nuclear retaliation.

There may have been some tendency on the part of Western diplomats to rely rather heavily on this facet of the situation. "What protects Berlin," Raymond Aron wrote in 1957, "is not the Soviet conviction that the Western powers would prefer death to the abandonment of the capital, but the doubt which subsists in spite of everything in the minds of the Soviet leaders about

²⁵ *New York Herald Tribune*, March 10, 1959. The earlier robbery is reported in *New York Herald Tribune*, February 20, 1959.

²⁶ Two weeks after this hypothesis was advanced in public, the management of the Chemical Corn Exchange system, as though they had reached a similar conclusion, began the installation of a system that would photograph customers *automatically*. In the course of the installation, a customer obtained several thousand dollars from a teller with a note that simply began, "Give me . . ." and made no mention whatever of weapon or punishment. He walked out with the money past workmen installing the new apparatus. (On a subsequent attempt, he was unmasked as an accountant who had taken, each time, a short break from his officework.)

the Western reaction."²⁷ That doubt may still shield Berlin. It is hard, at this moment,²⁸ to see what else does.

Nevertheless, Acheson points out in his book, such a policy has risks. "To put forward as a policy a threat which is incredible may be to stay within the yet uncharted bounds of permissible fraud in politics, but it is highly dangerous for the country."²⁹ This is quite true, as far as it goes. It may, in other words, be dangerous to be a sheep in wolves' clothing. But the risks are more widespread than that. It can also be dangerous to commit oneself to the most highly credible threats: if there is a significant chance that they will not be credible enough. And that there always will be.

Our threats can fail for a great variety of reasons. For example: (a) if we underestimated the "required credibility" (the opponent's critical risk) for this threat, that is, if we misread the opponent's payoffs; (b) if the opponent did not find our commitment as credible as we expected, credible enough; (c) if he should be irrational, impulsive, careless; (d) if he should be committed, himself, to the action we are trying to deter.

Deterrence has been perhaps too often described in terms of "posing risks." When it involves commitment, it also means *taking* risks: the risks that it will fail, that the threat will have to be carried out because the threatener has left himself no choice. The risks must be calculated; but how will the results be checked? Not only are the estimates subjective; they are estimates, largely, of subjective variables. What are the opponent's payoffs; what outcomes does he expect from given strategies, how does he evaluate

²⁷ Raymond Aron, *On War* (New York, 1959), p. 113.

²⁸ March 10, 1959: the date that this Lowell Lecture was delivered. At the time of this editing, July, 1959, there seems little reason to change the remark.

The Presidential press conference next day, March 11, 1959, threw further insight into United States strategy at that point in the Berlin crisis. The President made the following comments: "We are certainly not going to fight a ground war in Europe. . . . You wouldn't start the kind of ground war that would win in that region if that were going to make the way you had to enforce your will. You have got to go to other means." He said further: "I don't know how you could free anything with nuclear weapons. . . . Destruction is not a good police force. You don't throw hand grenades around the streets to police the streets so that people won't be molested by thugs. Well, now, this is exactly the way that you have to look at nuclear war, or any other. . . . And, I must say, to use that kind of a nuclear war as a general thing looks to me a self-defeating thing for all of us, because after all, with that kind of release of nuclear explosions around this world, of the numbers of hundreds, I don't know what it would do to the world, and particularly the northern hemisphere; and I don't think anybody else does. But I know it would be quite serious."

These views seemed to some listeners to be contradictory. The question was raised, that if ground war were ruled out and nuclear war didn't free anyone, what threat did the President have in mind? The President's reply: "I didn't say that nuclear war is a complete impossibility."

In the context of his other remarks, it was clear that the threat being posed was that of a *possibility* (which, in fact, the President tended to minimize) of a nuclear response. As Aron had suggested, this could be effective, given a sufficiently low critical risk in the opponent. But it might also be pointed out that if such a threat works today in deterrence, it might work tomorrow in blackmail.

²⁹ Acheson, *Power and Diplomacy*, p. 47.

those outcomes? To what extent is he a gambler? What alternatives does he see? What does he expect me to do; how does he see my payoffs, my expectations? How will my tactics influence his expectations? How likely is he to have committed himself? To act carelessly, irrationally? What risks will he take, and what does he think the risks are?

Our analysis tells us the relevance of those questions; it suggests concepts for representing and comparing the answers, some hints for drawing implications from the answers; but it does not give the answers. To find those answers is not within the scope of our logic, nor, as yet, of any science: it is an art.

Our policy of deterrence has been described, and defended, in these apt terms, written after Korea, Indochina, and the first shelling of Quemoy: "You have to take chances for peace, just as you must take chances in war. Some say we were brought to the verge of war. Of course we were brought to the verge of war. The ability to get to the verge without getting into the war is the necessary art."³⁰

But blackmailers too can calculate risks—and take them. They too can go to the verge of war; and this fact has an important bearing on the risks of deterrence.

In this abstract discussion, we have examined the anatomy of blackmail. In the next lecture, we shall hear the *sound* of blackmail; the words that Adolph Hitler spoke, and their echoes, that won him half of Europe before the firing of a shot. There is the artist to study, to learn what *can* be hoped for, what can be done with the threat of violence.

³⁰ John Foster Dulles, quoted in, "How Dulles Averted War," *Life*, Jan. 16, 1956, p. 78.