ALVIN I. GOLDMAN

TOWARD A THEORY OF SOCIAL POWER

The concept of power has long played a significant role in political thought, and recent decades have witnessed many attempts to analyze power and provide criteria for its measurement.[1] In spite of this impressive literature, however, our understanding of power remains inadequate. Specifically, no fully comprehensive conceptual framework exists within which questions about power can be formulated precisely and dealt with systematically. In the absence of such a framework it is difficult to investigate empirical questions, such as the extent to which a country is dominated by a 'power elite,' and it is hard to discuss normative issues, such as the relationship between power and freedom, or the relationship between equality of power and justice.

In this paper I shall outline a theory of power the ultimate aim of which is to shed light on empirical and normative questions in the political domain. To achieve this aim, however, it is wise to broaden our perspective. The domain of power is not confined to the political realm, narrowly conceived: employers have power over employees and teachers over students. Nor is power purely an inter-personal matter: men have power over nature as well as over other men. Although I am primarily interested in social and political power, I shall construe the notion of power broadly enough to cover all of these areas. There is an even wider sense of 'power' with which I shall not be concerned, the sense in which we speak of the power of an engine or a machine. I shall confine my attention to the sense in which a *man* or a *group* of men are said to have power. In this endeavor I shall be in the spirit of Hobbes, who wrote: "The power of a man, to take it universally, is his present means to obtain some future apparent good."

A theory of power must enable us to account for the fact that Nelson Rockefeller is, on the whole, an extremely powerful person. This does not imply that he has power over every man or every issue. There are many

issues over which Rockefeller lacks power completely, e.g., what grade a student receives in my course. Nevertheless, a person's total power is clearly related to his power over particular issues. Thus, in the first section of the paper, I shall explain what it is for an individual to have power with respect to a given issue, for example, who obtains a certain appointment, or whether it rains on Wednesday. Social theory does not restrict its interest to the power of individuals, however; of equal significance is the power of groups. In the second section, therefore, I shall turn to the nature of collective power. An analysis of collective power is needed to appraise the power of a nation, of the Pentagon, or of groups such as automobile manufacturers and ethnic minorities. Moreover, to the extent that a person's power is a function of his membership in groups that have collective power, an understanding of collective power is necessary to complete our account of the power of an individual. The first two sections of the paper primarily concern the conditions in which a person or group has at least *some* power on a given issue. A full-fledged theory, however, must enable us to make comparisons of power. While the Secretary of Defense has considerable power over military policy, the President has even more power. In the third and fourth sections, then, I shall address myself to the problem of comparisons of power and degrees of power. Finally, in the fifth section, I shall discuss the nature of power over a *person* (as opposed to an *issue*) and to the nature of *overall* power.

Throughout the paper I shall of course try to capture common-sense intuitions and judgments about power (ones that would be shared by people of different political persuasions). My primary concern, however, is not to canvass every use of 'power' in everyday speech, but rather to embark on the construction of a theory. Perhaps it is an inevitable feature of philosophical theorizing that certain intuitions and usages are emphasized while others receive merely passing attention. But the benefits of theory construction can make this price well worth paying.

I. INDIVIDUAL POWER

The central idea in the concept of power, I suggest, is connected with *getting what one wants*. An all-powerful being is a being whose every desire becomes reality. An all-powerful dictator is one whose every desire for state policy becomes the policy of the state. One man X has another

man Y 'in his power' if what happens to Y is a function of what X wants to happen to him. The central notion, in all of these cases, is that a powerful man is a man whose desires are actualized, i.e., a man who gets what he wants.

It would be wrong to conclude that whenever one gets what he wants one must have power in the matter. A farmer may want rain on a particular occasion and may happen to get what he wants, but it does not follow that he has power or control over the weather. Even a person who *regularly* gets what he wants need not be powerful. The Stoics, and Spinoza as well, recommended that one form one's desires to accord with what can realistically be expected to happen in any case; they regarded freedom as conformity of events with actual desires, or rather, as conformity of desires with events. But as an account of power this is inadequate. To take an extreme example, consider Robert Dahl's case of the 'chameleon' legislator who always correctly predicts beforehand what the legislature is going to decide, and then forms a desire or preference to accord with this outcome.[2] The chameleon always gets what he wants, but he is not one of the more powerful members of the assembly.

What these cases show – and what is probably clear from the outset – is that an analysis of power cannot simply concern itself with what an agent actually wants and actually gets, but must concern itself with what he *would* get on the assumption of various *hypothetical* desires. To say that S is powerful is not to say that he usually gets what he in fact wants, but that whatever he wanted he would get, no matter what he might happen to want. To explicate the concept of power, we must appeal therefore to subjunctive conditionals. Although we shall later discuss what it is for a person to be powerful in general, we begin by asking what it is for a person to have power over, or with respect to, a single issue or event. The following notation will be used. 'E' stands for an issue, e.g., whether or not it rains at a particular time and place; 'e' and 'not-e' stand for possible outcomes of that issue, e.g., the occurrence of rain and the nonoccurrence of rain, respectively. To assess the claim that person S has power with respect to (henceforth to be written: 'w.r.t.') an issue E, let us ask what would happen (a) if S wanted e to occur, (b) if S wanted not-e to occur, and (c) if S were neutral between e and not-e. If E is the issue of whether or not it rains (at a particular time and place), we would consider whether or not it would rain (a) if S wanted rain, (b) if S wanted non-

rain, and (c) if S were neutral between rain and non-rain. For any issue E, there are eight (logically) possible situations S might be in w.r.t. E, each represented by a function that maps attitudes of S vis-à-vis E onto outcomes of E. These eight situations are shown in the table below, in which the following notation is used: '$N(S, e)$' stands for 'S is neutral between e and not-e,' '$W(S, e)$' stands for 'S wants e', and '$W(S, -e)$' stands for 'S wants not-e'. The if-thens are subjunctive conditionals. Situation (4), for example, is a situation in which outcome e would occur if S were neutral on the issue, or if S wanted e to occur, but outcome not-e would occur if S wanted not-e to occur.

	If $N(S, e)$ then	If $W(S, e)$ then	If $W(S, -e)$ then
(1)	e	e	e
(2)	e	$-e$	e
(3)	e	$-e$	$-e$
(4)	e	e	$-e$
(5)	$-e$	$-e$	$-e$
(6)	$-e$	$-e$	e
(7)	$-e$	e	e
(8)	$-e$	e	$-e$

In two of these possible cases, viz. (1) and (5), S is impotent w.r.t. the issue. In (1) it would rain no matter how S might feel about it. In other words, S's attitude would make no difference to the outcome, though if he happens to have the 'right' desire, he will get what he wants. In cases (2) and (6) S is even worse off than impotent; he is '*counter-potent*'. In these possible cases it would rain if S wanted non-rain and it would not rain if S wanted rain. Cases (3) and (7) are rather anomalous, and I have no name for them. But, like the four previous cases, they are not ones in which S has power w.r.t. the rain. Turning finally to (4) and (8), we find cases in which S does have power w.r.t. the rain. In these cases S would get his way on the issue no matter which outcome he might prefer. A desire for rain would lead to rain and a desire for non-rain would lead to non-rain. The difference between (4) and (8) is that in the former, rain would occur if S were neutral on the issue, whereas in the latter, non-rain would occur if S were neutral. In both cases, however, a desire by S for either outcome would lead to the occurrence of that outcome.

The following analysis of power is suggested by the foregoing.

S has power w.r.t. issue E if and only if

(I)   (1) *If S wanted outcome e, then e would occur*, and
       (2) *If S wanted outcome not-e, then not-e would occur.*

This analysis, I believe, is very much on the right track. But the following difficulty presents itself. Suppose that S is totally paralyzed and incapable of action. Another person, S^*, has the ability to control the weather; but S does not know of S^*'s existence, and even if he did, he would have no idea how to communicate with him. Unbeknownst to S, however, S^* can detect S's desires by ESP (or by appropriate gadgetry attached to his central nervous system), and S^* has freely decided to make the weather conform with whatever desire S has vis-à-vis the weather. In this case S satisfies (I) w.r.t. the weather, but it is doubtful that we would credit S with power w.r.t. the weather. Although S would get what he wants vis-à-vis the weather, this is not because of anything *he* would *do*.

A refinement in our analysis seems to be called for, a refinement that incorporates the element of action. To formulate this refinement, let us introduce the notion of a *basic act-type*. There are certain types of acts, e.g., raising one's hand, taking a step, and uttering certain sounds, which have the following properties: (a) in ordinary circumstances, if a person wanted to perform such an act, he would perform it, and (b) his ability to perform it is independent of knowledge or information concerning other acts that would have to be performed in order to perform it.[3] In short, a basic act–type is one that a person can do 'at will', an act–type that is 'directly' under his control. Now frequently, in order to achieve a desired outcome, a person has to perform an appropriate sequence of basic acts, and these acts have to be performed at appropriate times. It is necessary, therefore, to attend to temporal matters. First, the issue in question must be clearly specified by indicating its time (i.e., the time of its outcomes). Who will be elected mayor in 1972 is a different issue from who will be elected mayor in 1974; S may have power w.r.t. one of these issues but not the other. Similarly, whether or not it will rain in Ann Arbor on Tuesday is a different issue from whether or not it will rain in Ann Arbor on Wednesday; S's having power w.r.t. one of these issues does not ensure his having power w.r.t. the other. We shall be concerned, then, with the issues that have 'built in' times.[4] This will be indicated by placing the

temporal reference in parentheses. Secondly, we need to specify the time
at which a person has power w.r.t. an issue. If, at t_1, there are sequences
of acts available to S that would lead to each of the outcomes of E (at t_{10}),
then S may have power, at t_1, w.r.t. E (at t_{10}). But if S fails to act appro-
priately between t_1 and t_5, he may no longer have power, at t_5, w.r.t.
E (at t_{10}). With these points in mind, we may propose the following anal-
ysis of (individual) power.

> *S has power, at t_1, w.r.t. issue E (at t_n) if and only if*
> (I')   (1) *There is a sequence of basic act-types such that*
>   (a) *if S wanted e, then he would perform these acts*
>   *at appropriate times between t_1 and t_n, and*
>   (b) *if S performed these acts at these appropriate*
>   *times, then e would occur (at t_n);*
>   (2) *There is a sequence of basic act-types such that*
>   (a) *if S wanted not-e, then he would perform these acts*
>   *at appropriate times between t_1 and t_n, and*
>   (b) *if S performed these acts at these appropriate*
>   *times, then not-e would occur (at t_n).*[5]

Let us see how (I') would apply to the rain example. The issue here is
whether or not rain occurs (at t_n). Let us assume that the state of the world
(in particular, the meteorological conditions) at t_1 is such that it is going
to rain at t_n unless S does certain things between t_1 and t_n. (I assume that
the occurrence or nonoccurrence of rain is *determined* by antecedent
events, including the acts of S. Determinism will be assumed throughout
the paper.) In order to prevent rain, S must disperse the clouds, or evap-
orate them, and this must be done before t_n. Now there are two possible
cases to consider.

Case 1: S is an ordinary fellow who has no way to disperse or evaporate
the clouds; in other words, no basic acts of his would lead to the disper-
sion or evaporation of the clouds prior to t_n. In this case, S does not
satisfy (I'). Of course S trivially satisfies the first conjunct of (I'), assum-
ing, at any rate, that he can perform some basic acts. For no matter what
sequence of basic acts he performs between t_1 and t_n, it *will* rain at t_n
(the rain will not be *caused* by any of his acts, however). But S does not
satisfy the second conjunct of (I'), for, *ex hypothesi*, no basic acts S could
perform would lead to non-rain at t_n. *Case 2*: S possesses a chemical that

evaporates rain clouds, and S has an airplane at his disposal that would enable him to spray the chemical onto the clouds (before t_n). Moreover, S knows that the chemical would evaporate the clouds, know how to fly the plane, and has other requisite pieces of information. In this case, S satisfies the second conjunct of (I') in addition to the first conjuct. There is a sequence of basic acts – viz., acts by which S would place the chemical into the plane, acts by which he would fly the plane into the atmosphere, and acts by which he would spray the chemical onto the clouds – such that (a) if S wanted non-rain he would perform this sequence of acts (at appropriate times), and (b) if S performed them, non-rain would occur (at t_n).

It must be stressed that the conditionals in (I') are subjunctive conditionals, not *causal* subjunctive conditionals. If they were construed as *causal* conditionals, then even in Case 2 S could not be credited with power w.r.t. E. This is because there is no sequence of basic acts S can perform that would *cause* the occurrence of rain at t_n (although there are many sequences of basic acts he can perform such that, if he performed them, rain would occur at t_n).

It might be argued that the use of causal subjunctives could simplify our analysis; instead of requiring that S be able to obtain *both* outcomes of E, we could merely require that S be able to *cause* (at least) *one* outcome of E. (As long as ordinary subjunctives are employed rather than causal subjunctives, the requirement that *two* outcomes be achievable must be retained; otherwise S would qualify for having power w.r.t. the rain even in Case 1. Using causal conditionals, however, it might be sufficient to require that S be able to cause *one* outcome.) Although reliance on causal conditionals probably would simplify the analysis, I believe it would also prove less illuminating. By relying on ordinary subjunctives only, we shall be forced to take a careful look at problems which might be neglected if we allowed ourselves the luxury of the causal idiom. Causal terminology is especially unhelpful in dealing with the kinds of cases of paramount interest to a theory of power, i.e., cases in which an outcome is a function of the actions of numerous agents. In thinking about distributions of power and degrees of power among many persons and many groups, the use of causal terminology is likely to obscure the crucial questions rather than illuminate them. It seems advisable, therefore, to avoid all reliance on the concept of causation from the outset.

A different possible objection to our analysis of power is that it seems to neglect the most important ingredients of power. In most of the literature on power, such things as wealth, authority, status and similar 'resources' play a crucial role; but they make no appearance whatever in (I'). Is this not a devastating omission? Now although such resources are not explicitly mentioned in our analysis, it should be clear that the satisfaction of (I') in any particular case depends on precisely such factors. Our analysis makes no reference to the existence of certain chemicals, or the availability of an airplane, but it is just such things that make it true that S has power w.r.t. the rain (at t_n). Similarly, although our analysans makes no explicit reference to institutional hierarchies or to positions of influence, it is just these sorts of things that determine whether a person has power w.r.t. the granting of a government contract. If S has control over a defense contract, this is not simply because he has a certain basic act repertoire; it must be because he occupies a position of authority in the governmental structure, or perhaps because he is in a position to influence officials through credible threats or offers. That these are the sorts of resources that give rise to power, however, does not imply that they ought to be mentioned explicitly in the analysis of power. Indeed, it would be foolish to try to construct an analysis that itemizes relevant resources. For even a single issue, the number and variety of potentially relevant resources is endless; and it is surely wholly impossible to say what resources are necessary and sufficient for issues *in general*. What is important, then, is that an analysans imply the *existence* of an appropriate set of resources, without necessarily characterizing these resources 'intrinsically'. This is what we accomplish in our analysans by the use of subjunctive conditionals.

Our analysis of power may be compared to a conditional analysis of disposition terms like 'soluble' or 'fragile'. In analyzing 'X is soluble in water' as 'If X were immersed in water, it would dissolve', we say nothing specific about the internal structure of X, the structure in virtue of which it is true that X would dissolve if immersed in water. Nevertheless, there must *be* some actual structure which makes this conditional true. Moreover, we may construe the hypothetical statement 'If X were immersed in water, it would dissolve' as asserting the existence of such a structure.[6] That is, we may construe it as asserting: 'The structure of X (and the structure of water) is such that if X were immersed in water, it would

dissolve'. The case of power is quite analogous. In saying that outcome e would occur if S were to perform certain acts, we do not indicate which facts about the world, which resources possessed by S, make this conditional true. Nevertheless, there must be some such facts or resources. Indeed, we may construe our subjunctive conditionals as asserting their existence. That is, we may construe the force of (1) (b) as follows: 'The state of S's resources (or the state of the world) is such that if S performed the indicated sequence of basic acts at appropriate times, then outcome e would occur (at t_n)'. We may say, therefore, that resources such as wealth, authority, reputation, attractiveness, friendship and physical location play the same sort of role vis-à-vis power as molecular structure plays vis-à-vis solubility. It is the possession of such resources that confers power, or, if you like, that *is* power.

The preceding remarks give us insight into the kinds of conditions that must obtain in order that certain outcomes would occur if S performed certain sequences of basic acts. In other words, we gain some insight into the conditions underlying clauses (1) (b) and (2) (b) of (I'). But what of (1) (a) and (2) (a)? What conditions must hold in order that S would select an appropriate sequence of basic acts if he wanted outcome e and in order that S would select an appropriate sequence of basic acts if he wanted outcome not-e? The answer is: S must have appropriate items of knowledge or belief. In order that S select appropriate sequences of basic acts – i.e., sequences that would really lead to the desired outcome – it is not sufficient that S simply *want* that outcome. Unless he has knowledge, or belief, concerning *which* acts would lead to the desired outcome, he might select acts which do not lead to it at all. We can imagine cases, indeed, in which there *are* sequences of basic acts which, if performed by S at appropriate times, would lead to whichever outcome he might desire, and yet where S is *counter-potent* w.r.t. the issue! S might be so confused or misinformed that if he wanted e he would perform acts that would lead to not-e, and if he wanted not-e he would perform acts that would lead to e. Under these circumstances, we would hardly say that S has power w.r.t. E. We may conclude, therefore, that the possession of power w.r.t. an issue depends not only on the possession of physical or social resources, but also on the possession of informational resources. Hence the maxim 'Knowledge is power'.

It must be acknowledged that there are certain uses of the term 'power'

– especially as a count-noun – in which informational resources seem irrelevant. To take an example of Rogers Albritton's, a man who is endowed with a capacity to strike people dead by uttering a magic formula can be said to possess this special *'power'* even though he does know that he has this power, and even though he has no idea what the formula is. Similarly, we may speak of an officer of an organization as having certain powers even if he happens to be ignorant of the fact that he has these powers. In this use of the term 'power', a power simply seems to be a resource of a crucial sort, though not necessarily a resource which is sufficient, by itself, to ensure any particular desired outcome.

Having noted this use, however, I propose henceforth to ignore it. I am interested primarily in the conditions under which a person has power over, or w.r.t., an issue, and the possession of power w.r.t. an issue *does* seem to require informational resources. Suppose S is standing in a large mansion which contains, unknown to S, a hidden button; if this button is pressed, New York City will be destroyed. Suppose, moreover, that there is a sequence of basic acts which, if performed by S, would uncover the button and place him in a position to destroy the city. Does it follow that S has power w.r.t. the destruction of the city (say, in the next ten minutes)? If S has no idea whatever where the button is, and if S has no way of finding out where the button is (in the next ten minutes), then I think it is clear that S does not have power w.r.t. this issue. Thus, the absence of relevant information implies the absence of power.[7]

A precise statement of the required information, however, is difficult to formulate. It is not necessary that S *know*, for each outcome e and not-e, which sequence of basic acts would be appropriate; it is not even necessary that he *believe* – in the sense of *believe it to be more probable than not* – of any sequence that it would lead to the desired outcome. It is sufficient for S to believe, of a certain sequence which happens to be appropriate, that it is *more likely* than any other sequence to lead to the given outcome; this is sufficient even if he thinks that the chances of its leading to that outcome are very small. To introduce some terminology to cover this possibility, we may say that it is sufficient that S *'epistemically* favors' that sequence as a means to that outcome.

Further complications are introduced by temporal considerations. In order for S to have power at t_1 it is not necessary that S be able *at* t_1 to select an entire sequence of basic acts. If, at t_1, S has a kit for assembling a

harpsichord, then he may have power, at t_1, w.r.t. the issue of whether or not there will be a harpsichord in his house at t_n (say, two months later). But there certainly is no entire sequence of basic acts which S epistemically favors, at t_1, as a means to obtaining the outcome of there being a harpsichord in his house at t_n. There are certain basic acts which S epistemically favors as *initial* members of such a sequence, viz., acts that would enable him to read the instruction manual. Moreover, once he performs these initial acts, his reading of the instruction manual will lead him to form beliefs concerning further basic acts to perform. But at no point does he have the entire sequence of appropriate basic acts in mind. Extrapolating from this case, we can say that S would obtain an outcome e, if he wanted it, as long as the following is true:

There is a sequence of basic acts $A_1, ..., A_i, ..., A_n$ such that

(A)   at t_1 S epistemically favors A_1 as the first act to perform as a means to e,

(B)   for every i $(1 \leqslant i \leqslant n\text{-}1)$, if S performed the first i members of the sequence from t_1 through t_i (assuming that each act is performed at a single moment), then S would epistemically favor, at t_{i+1}, act A_{i+1} as the next act to perform as a means to e, and

(C)   if S performed each of the acts A_i at time t_i $(1 \leqslant i \leqslant n)$, then e would occur (at t_n).

In the foregoing discussion I have talked as if the only factors that determine which basic acts S performs are, first, his desire for a particular outcome of E, and, second, his beliefs concerning the various means available to him to secure this outcome. In general, however, other desires and aversions come into play as well, not just S's attitude vis-à-vis the outcomes of E. For example, although S wants e to occur, and although he epistemically favors sequence $A*$ as a means to e, he may choose a different sequence of basic acts because he expects $A*$ to be very *costly*, that is, because he thinks $A*$ will have consequences he wants to avoid. In short, an agent's choice of action normally depends on more than one desire, and since more than one desire is involved, the relative strength of these desires is also an important factor. In the present section, however, I am intentionally neglecting these complications. I am assuming that S's desire for e (or for not-e) is the *only* motivating factor in his selection of a

course of action. Since it is the only motivating factor, its strength or inten-
sity is of no significance. If S wants e he will perform whatever acts he
epistemically favors as a means to e, and if he wants not-e he will perform
whatever acts he epistemically favors as a means to not-e. In Section IV I
shall drop these simplifying assumptions, but they will be retained until
then.

Since our analysis makes central use of subjunctive conditionals,
several remarks on their interpretation are in order. My approach to sub-
junctive conditionals follows the general lines of the analysis proposed by
Robert Stalnaker.[8] On this analysis, we assess the truth of any conditional
'If A then B' by considering a possible world in which A is true and which
otherwise differs minimally from the actual world. A conditional of this
form is true if and only if B is true in that possible world. Now if A is true
in the actual world, the possible world we select is the actual world; but if
A is contrary to fact, we must select some non-actual possible world.
The tricky matter here is to select the respects in which this possible world
should resemble the actual world and the respects (other than A itself) in
which it may differ. In other words, we must decide what, in addition to
A, is to be counterfactualized and what is to be held constant.[8a] One con-
straint is that the possible world must be a nomologically consistent world
(using the laws of nature of the actual world); but more must be said
about the selection of a possible world.

When we are interested in S's power, at t_1, w.r.t. E (at t_n), we begin by
counterfactualizing S's desire vis-à-vis E, more specifically, his desire
vis-à-vis E *from t_1 through t_n*. While counterfactualizing this desire,
however, the following three things are held constant: (1) the basic
act repertoire of S, (2) the set of beliefs that S has at t_1, and (3) S's re-
sources at t_1. The notion of resources is here construed very broadly, to
include not only physical conditions, such as the presence of clouds, but
also the acts and inclinations of other persons (*at t_1*). Thus, to say that we
hold constant the *resources* S has at t_1 is to say, roughly, that we hold
constant *the state of the entire world at t_1* – with the exception, of course,
of S's own desire vis-à-vis E and whatever is nomologically implied by that
desire (e.g., the state of his brain). Now once we have counterfactualized
S's desire vis-à-vis E, other counterfactualizations wil have to be made in
order to obtain a nomologically consistent world. In particular, if S's
desire from t_1 through t_n is different from his desire vis-à-vis E in the ac-

tual world, the basic acts he performs from t_1 through t_n will presumably be different (at least many of these acts). Moreover, if these acts are different, various other events that are causally connected with these acts will be different. Thus, once we have been forced to counterfactualize the basic acts S performs, we shall also have to counterfactualize numerous other events, including, perhaps, the acts of other agents and the beliefs S himself forms after t_1 (which in turn influence his subsequent acts). Of paramount importance is a possible change in the outcome of E. Since we are interested in S's power w.r.t. E, the crucial question is whether the counterfactual hypothesis that S desires, say, outcome e nomologically implies the performance of basic acts which nomologically imply the occurrence of outcome e (at t_n).

Our discussion of individual power has heretofore assumed that every issue has exactly two possible outcomes; and, working on this assumption, we have maintained that a person has power w.r.t. an issue only if he can obtain *each* of these possible outcomes. Both of these assumptions, however, are too restrictive. Most of the interesting issues in the social or political arena admit of more than two possible outcomes, and when this is so, it is not necessary, in order for a person to have power, that he be able to obtain *each* of the various possible outcomes. Suppose, for example, that money is to be allocated for a certain project, and there are 50 possible amounts that might be allocated: $ 1 million, $ 2 million, ..., $ 49 million, $ 50 million. Now suppose that although S is not able to ensure *whichever* of these 50 allocations he might desire, he is in a position to ensure any of the first *20* of these allocations. In this case, S clearly has power w.r.t. the issue of how much money is to be allocated for the project. On the other hand, if S is in a position to ensure any of the first *40* of these allocations, rather than any of the first *20*, then S has even more power w.r.t. this issue.

To accommodate this sort of case, we proceed as follows. First, instead of *defining* an issue in terms of a partition of outcomes, we think of a *single* issue as subject to a *variety* of different possible partitions. For example, we may take the single issue of *the weather* (at a particular time and place) and partition it into any of the following partitions, where each partition contains outcomes that are mutually exclusive and jointly exhaustive: (1) (a) rain, (b) non-rain; (2) (a) rain, (b) snow, (c) hail, (d) anything else; (3) (a) sunny, (b) cloudy, (c) precipitation; (4) (a) no precipitation, (b) less

than an inch of precipitation, (c) an inch or more of precipitation; etc.[9] Once we have the idea of different partitions of the same issue, we can make the following generalizations. If there is any partition of E into two or more mutually exclusive and jointly exhaustive outcomes such that S can obtain the occurrence of *at least two* of the outcomes of this partition, then S has at least *some* power w.r.t. E. Moreover, consider any partition P of E into possible outcomes $e_1, e_2, ..., e_n$, where $n \geqslant 2$. (For simplicity, I confine our attention to finite partitions.) Call the set of these n possible outcomes E^*. If we wish to compare S's power w.r.t. E in one possible world (where S possesses certain resources) with S's power w.r.t. E in another possible world (where he possesses different resources), we proceed as follows. We consider subsets E' and E'' of E^*, where E'' contains at least two members. Then if the following three conditions are satisfied, S has *more power* w.r.t. E in possible world W_2 than he has w.r.t. E in W_1: (a) in W_1 subset E' is the largest subset of E^* such that S can obtain whichever member he chooses, (b) in W_2 subset E'' is the largest subset of E^* such that S can obtain whichever member he chooses, and (c) E' is a proper subset of E''.[10]

II. COLLECTIVE POWER

I have focused until now on the power of a single person, but this barely touches the more important complexities in the topic of power. Most issues of interest in the social arena are issues in which many persons and many groups have some degree of power. Moreover, we are usually inclined to say, in such cases, that some of these people or groups have more power than others. Nothing I have said thus far, however, sheds light on these matters. I have said nothing concerning the power of groups of people, nor anything about comparisons of power between two or more persons (or groups). An adequate theory of power, obviously, must deal with these matters.

The problems that lurk in these areas can be conveniently introduced by a brief passage from an article in *New York Magazine* entitled "The Ten Most Powerful Men in New York." Having listed his choice for the ten most powerful men in New York, the author, Dick Schaap, writes as follows:

I offer only one theory in defense of the above list: if all ten men agreed upon the wisdom and necessity of a single, specific act affecting New York City, that act would take place, no matter how the rest of the city's eight million people felt.[11]

In passing, we may note two obvious deficiencies in Schaap's suggestion,

at least if it is regarded as a criterion for determining the ten most power-ful men. First, it is too strong. It is not necessary that the ten most power-ful men would *always* get their way despite the attitudes of *all* other New Yorkers. Secondly, the test does not ensure uniqueness. A number of different groups of ten men might each satisfy the test.

But let us reflect on other features of the test. It is noteworthy that Schaap's test does not require any one of the ten to be able to obtain out-comes *by himself*; that is, it does not require that there be any issue such that at least one of the ten would obtain different outcomes of that issue if he wanted them. The test can be satisfied even if each of the ten is only a *member* of a group whose *joint* preferences would determine the out-comes of issues. This seems perfectly appropriate if we are considering the power of the *group*. But by calling these men "the ten most powerful," Schaap also implies that *each* of them has considerable power. This raises an interesting question for the account of power given in Section I. If neither of two persons can *individually* get what he wants, but if the two of them can *jointly* get what they want, can either of them be credited *as a single person* with power w.r.t. the issue? This point requires an investiga-tion both into the nature of group or collective power and into the rela-tionship between collective power and individual power.

Another noteworthy feature of Schaap's 'theory' is the phrase 'no matter how the rest of the city's eight million people felt'. To complete Schaap's test of a proposed list of ten men, we must not only hypothesize that all ten agree in supporting a given outcome, but we must also hypoth-esize that all other New Yorkers oppose that outcome. Only if the ten men would achieve their outcome despite everyone else's opposition would it be true that they would achieve it "no matter how the rest of the city's eight million people felt," and only then would they qualify, on Schaap's test, as the ten most powerful. Schaap's requirement is reminiscent of a similar stipulation made by Max Weber, who defined 'power' as "the probability that one actor within a social relationship will be in a position to carry out his own will *despite resistance*...,"[12] and of C. Wright Mills' definition, "By the powerful we mean, of course, those who are able to realize their will, *even if others resist it*."[13] Our own analysis of power has made no reference to the resistance or opposition, either actual or hypothetical, of other persons. The place of this idea in an account of power must there-fore be explored.

In this section, however, we confine our attention to collective power. Suppose that you and I, both healthy and reasonably normal men, are standing behind a stalled Buick. If either of us alone pushes at it, the car will not budge; but if we both push simultaneously, it will move. Let E be the issue of the movement of the Buick (in the next several seconds) and let E be partitioned into two outcomes: (e) it moves, and (not-e) it does not move. If both of us want outcome e to occur, then we shall both push at the car and outcome e will take place. If both of us want not-e, neither of us will push and not-e will take place. Thus, if we jointly desire either outcome, that outcome will occur. This is a good reason to conclude that the two of us have *collective power* w.r.t. issue E. It appears, however, that neither of us has individual power w.r.t. E. True enough, if one of us wanted not-e, he would ensure that not-e would occur even if the other wanted e (or was neutral on the issue). But if only one of us wanted outcome e, while the other wanted not-e, the one who preferred e would not succeed in getting his way. It does not seem to be true of either of us, therefore, that for at least *two* outcomes of E, he would get his way on each outcome if he wanted that outcome. Actually, this conclusion is too hasty, as we shall see below. There is no doubt, however, that we must distinguish between collective power and individual power, and that collective power deserves study in its own right.

An analysis of collective power can easily be constructed along the lines of our analysis of individual power. For simplicity we confine our attention to two-outcome partitions.

A group of persons $S_1, ..., S_i, ..., S_m$ have collective power, at t_1, w.r.t. issue E (at t_n) if and only if

(II)   (1) *There is a set of sequences of basic act-types, a sequence for each person S_i, such that*

   (a) *if each person S_i wanted outcome e to occur, each would perform his respective sequence of acts at appropriate times between t_1 and t_n, and*

   (b) *if each person S_i performed his sequence at appropriate times, then e would occur (at t_n);*

(2) *There is a set of sequences of basic act-types, a sequence for each person S_i, such that*

   (a) *if each person S_i wanted outcome not-e to occur, each*

> *would perform his respective sequence of acts at appro-*
> *priate times between t_1 and t_n, and*
> (b) *if each person S_i performed his sequence at appropriate*
> *times, then not-e would occur (at t_n).*

Paralleling the case of individual power, there are two kinds or classes of resources that are relevant to the possession of collective power. First, informational resources are needed to satisfy clauses (1) (a) and (2) (a) of the analysis. Secondly, non-informational resources of various sorts are needed to satisfy clauses (1) (b) and (2) (b). Nothing especially distinctive is true of the class of non–informational resources, but some attention to informational (or epistemic) resources should be instructive.

In collective action toward a common goal, coordination is usually required in the selection of mutually supportive sequences of acts. For me to choose an appropriate sequence I may have to know what other members of the group are going to do; and similarly for each of them. Without information of this sort we may act discordantly despite good intentions. Often, therefore, acts performed in order to achieve a given outcome are designed to acquire information about future acts of partners in the undertaking. The nature of coordinative activity is a fascinating subject, which need not be expanded upon here.[14] It is (partly) the need for coordination, however, that makes the degree of organization or structural delineation of a group contribute to its power. An established pattern of division of labor facilitates the mutual selection of appropriate courses of action. The political power of lobbies and pressure groups, as opposed to that of random collections of individuals (e.g., until recently, consumers), is partly a function of this factor.

Two aspects of the problem of coordination should be distinguished. First, I may need information about the acts of others to decide *which* acts would be appropriate for me to perform. Secondly, even if I know which acts of mine are the ones most likely to lead to the desired outcome, I may need to know what others are going to do in order to assess *how likely* it is that the performance of these acts will be followed by the outcome. In particular, if I believe that others will not 'do their part', and if I believe that it will be very costly for me to perform my most appropriate sequence of acts, then even if I know which sequence of acts would be most appropriate for me to perform (as a means to the desired outcome),

I may choose not to perform them. Admittedly, this consideration intro-
duces the element of cost, which we resolved to abstract from until Sec-
tion IV. Nevertheless, it is sufficiently important in this context that it
should not be ignored.

Suppose that a small group of bandits are holding up a train containing
a large number of passengers. How shall we assess the collective power of
the passengers w.r.t. the issue of whether or not they will be robbed?
Suppose that the bandits 'have the drop' on the passengers, but that there
is a set of sequences of acts, a sequence for each passenger, such that if
they performed these sequences of acts, they would disarm the bandits
(with no harm to themselves) and foil the robbery. Assume further that
each passenger knows which acts would be the most appropriate ones for
him to perform as a means to foiling the robbery. This is not enough to
ensure that all would perform these acts if all wanted the robbery to be
foiled. The rub, of course, is that each passenger has little reason to believe
(indeed, has strong reason to disbelieve) that enough other passengers will
do their part. Since, for each passenger, it would be very costly if he did
his part (e.g., started to disarm the bandit nearest him) while few others
did theirs, each passenger would refrain from doing these acts, and the
robbery would succeed. A similar problem arises in assessing the power
of a large group of slaves over a small group of masters. If all the slaves
acted in unison, they would overwhelm their masters. But it does not
follow that they have much (or any) collective power over their masters.
Like the train passengers, the problem for the slaves is that each is insuf-
ficiently confident that rebellious action on his part would be supported
by others. There is an important respect, then, in which 'faith is power'.
To the extent that members of a group have greater confidence in the
reliability of their partners (and hence greater confidence in the efficacy of
their own acts as part of the larger group) the group itself has more
power, or is more likely to have at least *some* power w.r.t. a selected issue.

With a clearer grasp of the notion of collective power, let us next look
more closely at the relationship between collective and individual power.
In our discussion of the Buick example it was said that this is a case where
you and I have collective power w.r.t. its movement but where neither of
us has individual power w.r.t. its movement. Closer examination shows,
however, that this is not unconditionally true: it depends on further spec-
ification of the example. Suppose, first, that you are not going to push at

the Buick in the next few seconds, and that there are no acts I could per-
form that would induce you to push. In that case, there is no sequence of
basic acts I could perform that would lead to the car's moving (in the next
few seconds). Hence, it is correct to say that I lack individual power w.r.t.
the issue. But suppose now that, as a matter of fact, you *are* going to push
at the Buick in the next few seconds (and suppose that you have no incli-
nation *not* to push in case I push). In that case, (I') licenses us to say that I
do have individual power w.r.t. the issue, for your act of pushing serves
as a 'resource' of mine. If I wanted the car to move I would push at it,
and, given your pushing as a resource, this would lead to the car's moving.
If I wanted the car not to move, I would not push it, and in this case,
despite your pushing, it would not move. Thus, whichever outcome I
wanted would occur.

This example brings out the true contrast between collective power and
individual power (as we have defined these notions). The difference be-
tween them lies in the conditions that we counterfactualize in each case.
In making a judgment about the individual power of a person S at t_1, we
begin by counterfactualizing his desire (from t_1 through t_n) and his desire
only. The only other counterfactualizations that we allow are ones that
follow from, or are necessitated by, this initial counterfactualization.
In particular, we do not counterfactualize the desires or acts of other per-
sons unless those desires and acts would be affected by the difference in
S's desire. In making a judgment about the collective power at t_1 of a
group of persons, we begin differently. We begin by counterfactualizing
at the outset the desires of all of the members of the group. We ask what
would happen if they *all* wanted outcome e (from t_1 through t_n) and what
would happen if they *all* wanted outcome not-e (from t_1 through t_n).
Thus, in individual power we consider all persons other than S purely as
'resources' (or liabilities) of his; in collective power we consider all per-
sons outside the *group* purely as resources.

Let us return to the case where you and I have collective power w.r.t.
the movement of the Buick but neither of us has individual power (this
is the case where neither of us is *in fact* going to push, but where the car
would move if we both wanted it to move). According to the analysis of
Section I it should be concluded that neither of us has any power w.r.t.
the issue; for it was implied in Section I that a person has power w.r.t. an
issue only if he has *individual* power w.r.t. it. This assumption, however,

must now be called into question. As noted earlier, Schaap's formula implicitly denies it; for Schaap's formula allows a person to be one of the ten most powerful men in New York simply by being a member of a group that has collective power. What Schaap's formula suggests, then, is some sort of *distributive* principle: if a group of persons has collective power w.r.t. an issue, then every member of the group has power w.r.t. the issue. This formulation, however, is too strong. If group G has collective power w.r.t. E, another group G' can always be formed by adding to G some randomly selected, irrelevant person; and this new group, G', will also have collective power w.r.t. E. But we do not want to say of any randomly selected person that he has power w.r.t. E. A qualification is needed, therefore, to the effect that each member must be *non-dispensable* w.r.t. E. This notion can be explained as follows. If there is at least one outcome of E such that (a) if all members of G wanted that outcome it would occur, and (b) if all members of G except S wanted that outcome, whereas S opposed it, it would not occur, then S is a non-dispensable member of G w.r.t. E. In the Buick case, for example, you and I have collective power w.r.t. the issue, and each of us is a non-dispensable member of this group (w.r.t. this issue). We can now formulate the following principle.

> *A person S has some power w.r.t. issue E if there is some group G such that S is a non-dispensable member of G w.r.t. E and the members of G have collective power w.r.t. E.*

This principle proves useful in a variety of cases. In the Buick case, for example, although neither of us has individual power w.r.t. the issue, it seems plausible to say that each of us has some power, i.e., that neither of us is powerless w.r.t. it. The new principle licenses us to say this.

The principle also permits us to account for the power of each member of a legislature. To illustrate, consider an assembly of 100 legislators, in which 51 votes are required to pass a proposal and 50 to defeat it. Let E be the issue of whether a particular bill is passed and let us assume that the actual vote is 75 in favor and 25 opposed. Consider legislator S who in fact voted for the bill. Did he have power w.r.t. E? Assume that S did not have individual power, for the bill would have passed even if he had opposed it. Was S a non-dispensable member of a group which had collective power w.r.t. E? Yes. Let group G consist of 25 members of the assembly, including S, each of whom actually voted for the bill. Then, holding

constant the attitudes of the 25 original opponents of the bill and the remaining 50 supporters, we can say that if all members of G (including S) wanted the bill to be passed, it would have been passed, and if all members of G (including S) wanted the bill to be defeated, it would have been defeated. Thus, G had collective power w.r.t. E. But if all members of G *except for* S wanted the bill defeated, whereas S wanted it passed, then it would *not* have been defeated. Thus, S was a non-dispensable member of G w.r.t. E, and hence he had some power w.r.t. E.

Consider next a rather different case. Let E be the issue of whether Brown will be alive at noon today. At 11:00 a.m. Smith has a loaded gun in his hand, aimed at Brown, and Smith has the requisite beliefs such that he would kill Brown before noon if he wanted Brown dead and he would not kill Brown if he wanted him alive. A third man, Jones, is also in a position to kill Brown by noon. Moreover, Jones is resolved to kill Brown before noon if (and only if) Smith does not kill Brown. Finally, assume that Smith cannot influence Jones' action in this matter. Now in this situation we are surely inclined to say that Smith has some power w.r.t. E. According to the account of individual power, though, Smith does not have individual power (at 11:00) w.r.t. E.[15] Given the facts concerning Jones, it turns out that there is nothing Smith can do that would lead to Brown's being alive at noon; so Smith is impotent w.r.t. E, for Brown will be dead at noon no matter what Smith wants or does about it.

Using the new principle, however, we can account for the intuition that Smith does have at least some power w.r.t. E. For Smith is a non-dispensable member of a group (viz., Smith and Jones) which has collective power w.r.t. E. Smith and Jones have collective power w.r.t. E because if they both wanted Brown dead at noon that outcome would ensue, and if they both wanted Brown to be alive at noon they would perform acts leading to that outcome. Smith is a non-dispensable member of this group w.r.t. E because if Jones wanted Brown to be alive while Smith opposed it, then Brown would not be alive (at noon).

III. CONFLICT AND COMPARISONS OF POWER

We turn now to the significance of conflict or opposition in the criteria of power offered by Schaap, Weber and Mills. Schaap, it will be recalled, characterized the ten most powerful men in New York as the ten that

would get their way no matter how the rest of the city's people felt about it. Whether or not this criterion is adequate, its rationale is clear. To determine relative amounts of power between people (or groups), it is appropriate to ask who would 'get his way', i.e., who would get his preferred outcome, in case of conflict. Schaap does not suggest (nor do Weber and Mills) that power is present only in situations of actual conflict or opposition. What is suggested, though, is that comparisons of power can be made by ascertaining what would happen *if* there were conflict or opposition. Since comparisons of power are of central concern, let us reflect on on this matter. I shall not try to deal with the relative power of groups, nor with the relative power of three or more individuals. I shall confine my attention to comparisons between two individuals (although third-parties will have to be mentioned to clarify the nature of the two-person case).

Let us begin with an example. Jones and I are both standing next to an open door. Jones is a muscular 250-pounder and I a 145-pounder. Let E be the issue of whether or not the door remains open. If both Jones and I can rely on raw strength alone, it is pretty clear that he has more power than I do w.r.t. E. For there is a sequence of basic acts Jones can perform to ensure, no matter what basic acts I perform, that the door will stay open; and there is a sequence of acts Jones can perform that would ensure that the door will be closed, no matter what basic acts I perform. (Here we rely on the fact that Jones' basic act repertoire exceeds mine: he can exert a greater amount of pressure on the door than I can.) Thus, assuming Jones has the requisite beliefs, the door would stay open if Jones wanted it open, no matter how I felt about it; and the door would be closed if Jones wanted it closed, no matter how I might feel about it. Should Jones and I have opposing preferences, then, Jones' preferred outcome would be the one to occur.

In this example, Jones can perform a certain sequence of basic acts such that, no matter what basic acts I perform, the door will stay open; and similarly for the door being closed. Is this sort of relationship generally necessary in order that one person have more power than another w.r.t. an issue? Restricting our attention to partitions with two outcomes, this would be generalized as follows: S_1 has more power than S_2 w.r.t. E only if, for each of the two outcomes, there exists a sequence of basic acts S_1 can perform such that, for *any* sequence of basic acts S_2 might perform, the performance of the sequence by S_1 would lead to this outcome. Such

a requirement is clearly too strong. It demands the existence of a single course of action for S_1 that would 'win' in the face of all possible responses from S_2. A weaker requirement is therefore needed. The condition that naturally suggests itself next is this: For each of the two outcomes, no matter what basic acts S_2 might perform, there is a sequence of basic acts S_1 could perform that would lead to this outcome. This condition is *too weak*. The statement that for any course of action by S_2 there is a 'winning' course of action for S_1 is compatible with the statement that for any course of action by S_1 there is a 'winning' course of action for S_2.[16] So this requirement would not provide even prima facie grounds for thinking that S_1 has more power than S_2.

The next natural suggestion is to appeal to the game-theoretic notion of a winning strategy. What would be required is that for each outcome, e and not-e, there be a strategy or function F which assigns to S_1 an initial move and which, for every set of possible moves of S_2, assigns to S_1 a succeeding move (or moves), such that if S_1 were to abide by this function, that outcome would occur. In fact, however, even this requirement is too strong. The existence of a winning strategy for S_1 implies that S_1 can win no matter what S_2 does, in other words, even if S_2 adopts the best strategy available to him. But suppose that S_2 does not have information that would lead him to adopt (what is in fact) his best strategy. Then S_1 may be in a position to 'beat' S_2 even if S_1 lacks a *winning* strategy (i.e., a strategy that would win for S_1 against *all* possible strategies of S_2). In fact, S_2 might even have a winning strategy against S_1; yet if S_2 does not know what this strategy is, and has no way to find out, then S_1 may have more power than S_2 in spite of this. His having more power would simply consist in the fact that he would get his preferred outcome (by an appropriate course of action) if the two of them had opposing preferences.

Let us abandon the attempt to specify in detail what combinations of strategy-and-information would be necessary and sufficient for S_1 to have more power than S_2. Instead, we can give a simple analysis of this notion that parallels our earlier analyses of individual power and collective power.

> *At t_1 S_1 has more power than S_2 w.r.t. issue E (at t_n) if and only if:*

(III)   (1) *There is a pair of sequences of basic act-types, Σ_1 and Σ_2,*
             such that
             (a) *if S_1 wanted outcome e and if S_2 wanted outcome not-e,*
                 then S_1 would perform Σ_1 at certain times between t_1
                 and t_n and S_2 would perform Σ_2 at certain times between
                 t_1 and t_n, and
             (b) *if S_1 were to perform Σ_1 at these times and if S_2 were to*
                 perform Σ_2 at these times, then outcome e would occur
                 (at t_n);
             (2) *There is a pair of sequences of basic act-types, Σ_1' and Σ_2',*
             such that
             (a) *if S_1 wanted outcome not-e and if S_2 wanted outcome e,*
                 then S_1 would perform Σ_1' at certain times between t_1
                 and t_n and S_2 would perform Σ_2' at certain times between
                 t_1 and t_n, and
             (b) *if S_1 were to perform Σ_1' at these times and if S_2 were*
                 to perform Σ_2' at these times, then outcome not–e would
                 occur (at t_n).

We might wish to add clauses to the analysis to ensure that S_1 would get
his preferred outcome if S_2 were neutral on the issue, but this refinement
may be neglected. A far more important refinement concerns the elabora-
tion of the analysis to cover partitions of three or more outcomes.[17] But
this complication will also have to be omitted here.

    As in the case of collective action, selections of sequences of acts by
actors with opposing interests are typically interdependent. If S_1 and S_2
realize that they have conflicting preferences, each will be guided in his
choice of action by whatever information he has about the acts his oppo-
nent has performed or plans to perform in the future. As before, informa-
tional resources are crucial in determining the acts one would select, and
therefore important in determining one's relative power.

    In our door example we imagined that both Jones and I rely on our own
strength alone. In characteristic situations in the political sphere, however,
relative power depends on other assets, including positions of influence
over other persons that stem from authority, kinship, personal magnetism,
or other relationships. In seeking to achieve one's ends in opposition to
others, one frequently performs acts designed to call forth aid from other

persons. These acts might be orders, commands, or simple requests for help. To illustrate the importance of this, let us revise the door example to include a brawny and faithful companion of mine who is always willing and able to assist me. In this amended case it is no longer true that Jones would get his preferred outcome if his preference were in opposition to mine. If I wanted the door to be closed, for example, I would perform basic acts by which I would ask my friend for help, and then I would perform acts of pushing at the door which, in unison with my friend's pushing, would ensure that the door be closed even if Jones tries to keep it open. Thus, given my brawny and faithful companion, I would get whichever outcome I would prefer on the door issue, no matter how Jones might feel about it.

Is it true, however, that I have more power than Jones w.r.t. E? There may be a strong temptation to deny this. It might well be argued that it is only the power of the *two* of us, my companion and me, that exceeds the power of Jones. This view can be defended by pointing out that if we are considering *hypothetical* desires, and *hypothetical* sequences of acts, we should also consider hypothetical desires and acts by the companion. If we do this, we shall quickly see that I do not have the ability, given *any* desires and acts by my companion, to beat Jones at whatever he tries to do. I could be said to have this ability only on the assumption that the companion's desires and acts are wholly contingent on mine. But why should this contingency or dependency be assumed? Admittedly, I introduced the example by saying that the companion is willing to do my bidding, and it would violate this stipulation to hypothesize that I ask him for help but that he does not desire (on the whole) to do what I ask him to do. But why should such a hypothesis not be permitted? In talking about power we are already in the business of making counterfactual assumptions about various people's desires and acts. Why not make counterfactual assumptions about the companion's desires and acts as well? The problem, in short, is this: When making comparisons of power between individuals (or, indeed, in making non-comparative judgments of power), what should be regarded as fixed and what should be regarded as a candidate for counterfactualization?

The answer I propose is that it all depends on what sorts of power comparisons we wish to make. If we wish to compare Jones and myself, then *our* desires and acts are subject to counterfactualization, while the desires

and acts of everyone else are to be taken as contingent on ours – at least, if there is anything in the actual world that makes their desires and acts contingent on ours. In other words, in comparing the power of Jones and myself, the *initial* changes we make in constructing different possible worlds concern the desires of Jones and myself; the only other changes that are made are ones nomologically required by these initial changes. On the other hand, if we wish to compare the power of Jones, myself, and my companion, then we would make *initial* counterfactualizations involving my companion's desires as well. Because of these different counterfactuals, of course, different assessments of power are going to be made. There is nothing inconsistent or paradoxical about this, however; it merely reflects the fact that whenever judgments about power are made, certain things are subjected to counterfactualization and others held fixed. A judgment about a person's power that makes certain counterfactualizations cannot be expected to be identical with a judgment that makes different (initial) counterfactualizations. Within certain limits, however, a number of different counterfactual assumptions may legitimately be made.[18]

To illustrate these points consider a departmental secretary who is given substantial responsibility and initiative by the chairman. If we ask about the secretary's power w.r.t. a variety of issues (e.g., assignment of offices, teaching hours, classrooms, etc.) it is not so clear what should be said. On the one hand we may hold fixed, or constant, the chairman's propensity to go along with the secretary's decisions; in other words, we regard the chairman's trust as one of the secretary's *resources*. If so, we would attribute considerable power to the secretary. On the other hand, we may also compare the secretary's power with that of the chairman; we might ask what would happen if the secretary wanted one outcome and the chairman preferred a contrary one. Normally the chairman would win in such cases (let us suppose), and hence the secretary's power is less than the chairman's.

I think we can see from this that a clear, overall picture of someone's power (w.r.t. a given issue) demands a consideration of more than his *individual* power (as defined in Section I); it also requires attention to his power *relative to* other persons (and groups of persons). This is a point that is implicit in Weber's and Mills' characterizations of power as the ability to get one's way *despite* (possible) *resistance*. That individual power

by itself does not give us the whole story is evident in the following example. Suppose that S^* has a rain-making machine, and this gives him power w.r.t. the noon-time weather (since, in the natural course of events, it would not rain at noon). Suppose, in addition, that if S were humbly to beg and plead with S^* to make rain for him (S), then S^* would accede to S's wishes. If so, then according to our analysis of 'individual power', S has individual power w.r.t. the noon-time weather. But clearly, if S's only way of affecting the weather is by throwing himself on S^*'s mercy, we would not be inclined to credit S with much power w.r.t. the weather. This can be explained in terms of our theory by appeal to a comparison of power between S and S^*. It is obvious that S^*'s power w.r.t. the weather exceeds that of S. And this must be taken seriously in one's overall appraisal of S's power w.r.t. the weather.

An important problem arises here, however. Suppose that S is able to influence S^* to use his rain-making machine, not by *pleading* with S^*, but by *threatening* him with dire consequences should he (S^*) refuse to employ the machine to make rain. Here we would say that S has quite a lot of power w.r.t. the weather, more, at any rate, than in the previous example. This difference, however, cannot be captured by our account of comparative power, at least not by the account as it stands now. For even in this case our account of comparative power would have us say that S^* has more power than S w.r.t. the weather. This is because *if S^** wanted non-rain *on the whole* – despite S's threat – then non-rain would occur.[19] What this shows is that our analysis of comparative power must be supplemented or refined in some further way. As our discussion intimates, the needed refinement concerns the element of *cost*. The difference between pleading and threatening is that the latter imposes (prospective) *costs* on S^* which the former does not (at least not such heavy costs). By making use of the notion of cost, we can hope to develop an account of comparative power that will handle the difference we have detected.

IV. COST AND DEGREES OF POWER

The necessity of incorporating the element of cost into our theory was acknowledged early in the paper but posponed until now. Heretofore we have assumed that the only factor motivating the choice of a person's action is his desire or aversion for outcomes of issue E. As noted earlier,

though, this assumption is unrealistic. The performance of a sequence of acts often leads to undesired consequences; at a minimum, it typically involves the expenditure of valuable resources or assets. The performance of a sequence of acts, then, involves some *cost*. Now although a person may want outcome *e*, and although he may believe that the performance of sequence Σ is the only way for him to get *e*, he may decide not to perform Σ if the expected cost of performing it exceeds the benefit he would get from *e*. A directly related point we have neglected is the fact that a person can want an outcome to a greater or lesser degree. We have assumed until now that a person either wants *e*, wants not-*e*, or is neutral between them; we have not worried about the strength (actual or hypothetical) of a desire. Once the element of cost is introduced, however, strengths of desire and aversion must be included as well, for the choice of courses of action will depend on whether or not the value of an outcome *exceeds* the (expected) cost of obtaining it.

Introduction of the element of cost has immediate bearing on power. As we have seen, the expected cost of a sequence of action may dissuade a person from performing that sequence even if he wants the outcome to which (he believes) it will lead. It follows, therefore, that the cost of a sequence is a determining factor of whether or not a person *would* get a certain outcome *if* he wanted it (to such-and-such a degree). Hence, it is a determining factor of his power w.r.t. an issue.

The importance of cost in the analysis of power was first stressed by John C. Harsanyi.[20] He illustrates the idea as follows.

It is misleading to say that two political candidates have the same power over two comparable constituencies if one needs much more electioneering effort and expenditure to achieve a given majority, even if in the end both achieve the same majorities; or that one can achieve favorable treatment by city officials only at the price of large donations to party funds, while the other can get the same favorable treatment just for the asking.

Harsanyi stresses how inaccurate an analysis of power can be if it disregards the costs to an agent.

For instance, suppose that an army commander becomes a prisoner of enemy troops, who try to force him at gun point to give a radio order to his army units to withdraw from a certain area. He may well have the power to give a contrary order, both in the sense of having the physical ability to do so and in the sense of there being a very good chance of his order being actually obeyed by his army units – but he can use this power only at the cost of his life.... [I]t would clearly be very misleading in this situation to call him a powerful individual in the same sense as before his capture.

To generalize this point, we can say that if it is extremely costly for an agent to obtain certain outcomes of an issue – e.g., if it would require the sacrifice of his life, health or fortune – then he cannot be said to have much power w.r.t. that issue. The amount of power an agent has w.r.t. an issue is thus inversely proportional to the cost of obtaining outcomes of that issue.

The element of cost, then, enters into the analysis of power at the level of individual power. But it is even more important when we turn to comparisons of power, that is, when we consider what would happen in situations of conflict. When two persons have opposing preferences concerning an issue, one of them (or both) frequently tries to get his preferred outcome by the use of threats or sanctions designed to *deter* his opponent from performing certain acts. The intended effect of a threat is not to make the opponent literally *unable* to perform a certain sequence of basic acts; rather it is to make that sequence more *costly* for him to perform. To the extent that the ability to deter people by successful threats is an essential ingredient in the possession of power, it is vital that we incorporate the element of cost into our theory. How, exactly, may this element be incorporated?

The first problem is how to construe the notion of cost and how to measure it. The most promising approach, I think, is suggested by the economist's notion of opportunity cost. To say that the achievement of outcome e would be costly for S is to say that the activity by which S might obtain e would have consequences that are less desirable, or have lower utility, than an alternative course of action open to S. If obtaining outcome e would require an expenditure of money, then the cost of obtaining e is a function of the utility that would have accrued to S from using the money in some alternative way. If e can only be achieved at the cost of imprisonment, then the cost of obtaining e is a function of the difference in utility between going to prison and remaining a free man. It is evident that in determing the cost of an activity we must refer to the agent's desires, or utility assignments, for various outcomes. In this context, though, we consider the *actual* desires of the agent, not hypothetical desires. True, we continue to regard desires vis-à-vis the outcomes of the issue in question (e and not-e) as subject to counterfactualization. But other desires are held constant (at least so far as this is compatible with whatever counterfactualization is made).[21] Assume that a Senator

can obtain passage of a certain bill only if he engages in activity that would cost him his reelection. If he is in fact highly averse to losing the election, then the cost would be very high, and we shall say that his power w.r.t. the passage of the bill is correspondingly reduced. But if he in fact cares very little about reelection, then his cost is not so great, and he has more power w.r.t. the passage of the bill.

Problems of cost become complicated when we turn to comparisons of power, and hence to (potential) conflict situations. The complications here are twofold. First, both agents, S_1 and S_2, may be in a position to threaten his opponent with certain penalties, thereby imposing costs on the opponent. Secondly, however, the activity of posing these threats, or of making them credible, may be costly to the threatener himself. In assessing the relative power of S_1 and S_2, therefore, all of the following must be combined: (1) the costs that S_1 can impose on possible courses of action by S_2, (2) the costs *to* S_1 of imposing these costs on S_2, (3) the costs that S_2 can impose on possible courses of action by S_1, and (4) the costs *to* S_2 of imposing these costs on S_1.[22]

To combine these elements and incorporate them into our theory of power, let us make use of matrices resembling those of game theory. In constructing such matrices the following assumptions will be adopted. (1) Issue E is partitioned, as usual, into two outcomes, e and not-e. (2) Two agents, Row and Column, have various sequences of basic acts open to them. For simplicity we depict only three sequences of acts for each agent. (3) It is assumed that inter-personal assignments of utility can be made. (4) A three-by-three matrix, like Matrix 1 below, gives information on three subjects. First, it says what outcome would occur if Row and Column were to choose certain sequence of acts. This outcome is listed in the center of each box of the matrix. Matrix 1 says, for example, that if Row were to choose r_2 and Column were to choose c_3, then outcome e would occur. Secondly, the entry in the lower left-hand corner of each box indicates the cost to Row of the corresponding pair of sequences of acts. Thirdly, the entry in the upper right-hand corner of each box indicates the cost to Column of that pair of courses of action. (5) To ascertain the cost of a pair of activities to a given agent, say Row, we proceed as follows. We begin by *ignoring* all utilities assigned by Row to the outcomes of E. We next consider the (expected) consequences *apart* from e and not-e themselves, of the various alternative sequences of acts

Column

	c_1	c_2	c_3
r_1	0 e 0	-10 not-e 0	-30 e 0
r_2	0 e -1000	-20 e -1200	-40 e -1000
r_3	0 not-e -600	-10 not-e -500	-30 not-e -600

Row is labeled to the left of r_2.

Matrix 1

open to Row. These consequences might include the expenditure of a certain amount of money and energy, or it might include the commitment of a certain item of patronage, or it might include the undergoing of certain penalties or sanctions.[23] We then select, among all the alternatives open to Row, the sequence(s) of acts that would yield Row the highest utility. This sequence of acts has *zero* (opportunity) cost, and hence the numeral zero is inserted as the entry for Row in each box corresponding to that sequence of action.[24] To determine the cost entry for Row in all other boxes of the matrix, we compare the utility that would accrue to him from each combination of courses of action by Row and Column with the (zero) utility of his best alternative. This *difference* in utility is the cost to Row of that combination of activities, and it is inserted in his corner as a negative quantity. (6) Finally, it is assumed that each agent knows what the outcomes and costs would be (both to himself and to his opponent) of all possible combinations of sequences of acts.

To illustrate this procedure, consider Matrix 1. Looking first at the entries in the center of the boxes, it is clear that Row has a sequence of acts available to him – viz., r_2 – that would ensure the occurrence of e, and he has a course of action available to him – viz., r_3 – that would ensure not-e. It is clear, however, that both r_2 and r_3 would be very costly for Row to perform. Setting aside the potential benefit of outcome e or outcome not-e, the best sequence of acts open to Row – that is, the sequence that would yield the greatest utility (or the least disutility) – is r_1. In comparison with r_1, sequences r_2 and r_3 are very unattractive. But r_1 would neither ensure outcome e nor ensure outcome not-e. Hence

although there *are* sequences of acts available to Row that would ensure either of the two outcomes, these sequences of acts are very costly, and it is not clear that Row would perform either of these sequences even if he wanted the outcome it could ensure.

Assuming that Row is a 'rational' agent, what sequence of acts will he perform? That depends on whether he prefers e or prefers not-e, and it depends on how much, or how strongly, he prefers one to the other. It also depends, of course, on what Column is likely to do, which in turn depends on Column's preference as between e and not-e. Since we are interested primarily in comparisons of power, let us suppose that Row and Column have *opposing* preferences. This supposition can be satisfied in two ways: (a) Row prefers e while Column prefers not-e, and (b) Row prefers not-e whereas Column prefers e. Next, some assumption must be made concerning the *degrees* of preference. Let us assume, in each case that the degree of preference for *both* agents is the same, viz. 100 utils. We have, then, two (hypothetical) cases on which to focus: (A) Row prefers e to not-e by 100 utils while Column prefers not-e to e by 100 utils, and (B) Row prefers not-e to e by 100 utils while Column prefers e to not-e by 100 utils. What would happen in each case?

Only a little reflection is needed to see that, in each situation, Row and Column would perform sequences of acts that lead to *Column*'s getting his preferred outcome. In each case Row's degree of preference for a given outcome (whether e or not-e) would not suffice to motivate him to perform either r_2 or r_3. The cost of ensuring either outcome would outweigh the benefit from that outcome. Hence, Row would opt for r_1, thereby leaving it up to Column's choice of action to determine whether e or not-e occurs. Now in case (A), where we imagine Column to prefer not-e, he would select sequence c_2. Admittedly, sequence c_2 is more costly, by 10 utils, than sequence c_1. But Column is assumed to prefer not-e to e by 100 utils, and so he would be better off, all things considered, to perform c_2. By selecting c_2, given that Row selects r_1, Column gets his preferred outcome (not-e). By similar reasoning (indeed, even more obviously), Column would get his preferred outcome in case (B); for in this case Row would perform r_1 and Column would perform c_1, a combination that would result in outcome e. Thus, in both hypothetical cases, Column would get his preferred outcome; and this is true despite the

fact that Row has 'winning' strategies available to him, and knows that he has these strategies, for both e and not-e.

Let us step back now and take a new look at where our analysis leaves us. According to our initial conception, the idea of power was understood in terms of functional dependencies between the desires of an agent and the outcomes of an issue. We have stressed, however, that the outcomes of an issue are a function, not of desires *simpliciter*, but of *degrees* of desire. If Row's preference for e over not-e were not of the magnitude 100, but rather of the magnitude 2000, he would perform a different sequence of acts, and the outcome of the issue would be different. Now in the first three sections of the paper, we have talked about desiring e or not-e (or of being neutral between them). But the introduction of degrees of desire can help us add further refinements to our theory of power; specifically, it enables us to work toward *an* account of *degrees* of *power*.[25]

How are degrees of power related to functional dependencies between degrees of desire and outcomes of an issue? The answer is straightforward: *The extent of a person's power w.r.t. an issue is (ceteris paribus) inversely related to the degree of desire required for him to obtain a preferred outcome.* This answer is directly tied to considerations of cost. If obtaining a preferred outcome is very costly for an agent, then it will require a higher degree of desire to motivate him to get that outcome. (Cf. the captured commander.) And the greater the cost of getting an outcome, the less is a man's power w.r.t. that issue.

There is a clear, intuitive idea here that we seek to capture. If a person is in a position vis-à-vis an issue such that even the slightest concern, the merest whim, would suffice for him to get his preferred outcome, then he has great power w.r.t. that issue. But if a man is in a position such that only very great concern would make it worth the trouble, effort, or expenditure of resources to obtain his preferred outcome, then he does not have so much power w.r.t. the issue. If a mere whim on the part of a well-placed corporation executive would suffice to get an appropriate piece of legislation enacted, then he has great power w.r.t. that issue. I might be able to get a similar piece of legislation enacted, but only if I wanted it badly enough to go to a great deal of trouble and effort. This implies that my power w.r.t. this kind of issue is much smaller than the executive's. Similarly, if the executive and I are on opposing sides of

some issue, and if it is true that a moderate desire on his part for his most preferred outcome and a large desire on my part for my most preferred outcome would result in a victory for him, then his power exceeds mine. If we apply this idea in the case of Matrix 1, it is not at all clear that Row's power exceeds Column's.

It is a consequence of the current suggestion that sometimes a man who has less power w.r.t. an issue will get his preferred outcome over one who has more power on that issue. This can happen when the degree of preference of the less powerful man far exceeds the degree of preference of the more powerful one. Suppose, for example, that Row prefers outcome e to not-e by 2000 utils, while Column prefers not-e to e by only 100 utils. (Actually, it doesn't matter here what the strength of Column's preference is.) Then Row will choose r_2 and will get his preferred outcome. This does not conclusively falsify the suggestion that Column has more power on the issue than Row. There is no reason to suppose that a more powerful person on a given issue will always get his way over a less powerful one. On the contrary, it seems reasonable to define relative power in such a way that a powerful but relatively uncaring man may lose out to a less powerful but strongly motivated fellow. What a definition of relative power must ensure is that a man who is powerful w.r.t. a given issue will get his way if the issue is one of considerable importance to him, at least if he is contending against someone whose concern for the issue is comparatively small.

The foregoing discussion suggests that we may conceive of a two-person power relationship as a function which maps pairs of degrees of desire (of the two agents) onto outcomes of the issue in question. This may be illustrated with the aid of simple graphs. Let X and Y be two agents and let e and not-e be two outcomes of an issue. X's preference w.r.t. these outcomes is measured along the X-axis of a coordinate system. A positive number along the X-axis represents a degree of preference by X for e over not–e, and a negative number represents a degree to which X prefers not-e over e. The zero-point signifies indifference between e and not-e. Similarly, Y's preference as between e and not-e is measured along the Y-axis. A point in the coordinate system will then represent a pair of degrees of preference (or desire) by X and Y. The point $(+100, -200)$, for example, represents a preference of 100 utils for e over not-e on the part of X and a preference of 200 utils for not-e over e on the part of Y.
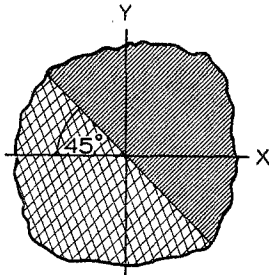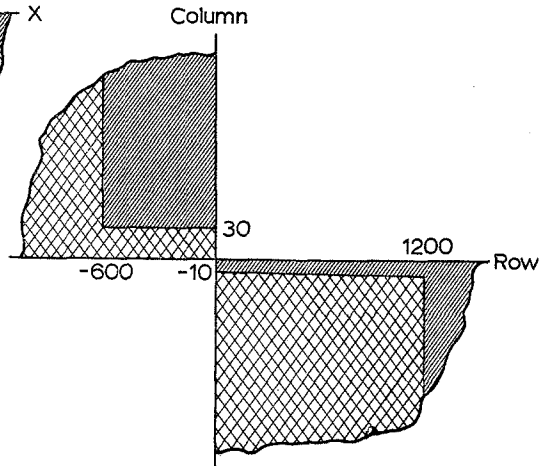
Graph 1
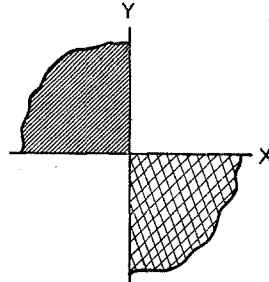
Graph 2
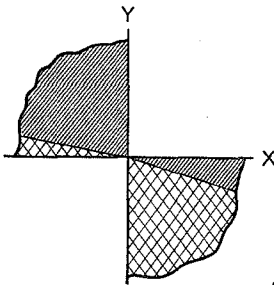
Graph 3

Graph 4

Now suppose that the power position of X and Y is such that if X were to want e to degree 100 and Y were to want not-e to degree 200 then e would occur. We shall represent this fact by darkening, or filling in, the point $(+100, -200)$. (Of course we cannot literally darken *points*, but appropriate *areas* can be darkened.) If, on the other hand, such a combination of desires by X and Y would result in not-e, then the point $(+100, -200)$ will be marked by cross-hatching. (Again, it is only areas that can be so

marked.) Using these conventions, we can proceed to represent various possible relationships of power of X and Y w.r.t. issue E. Indeed, we can give information concerning both their *collective* power on the issue, and their *relative* power.

Consider Graph 1 on page 255. Looking at the northeast and southwest quadrants, we see that X and Y have the highest degree of *collective* power w.r.t. the issue (at least vis-à-vis this partition). For if X and Y both prefer e to not-e then e will occur – no matter how small their degree of preference, and if they both prefer not-e to e then not-e will occur – again no matter how small their degree of preference. Looking at the northwest and southeast quadrants, we see what would happen if X and Y had opposing preferences. The graph tells us that, in case of opposing preference, the agent with the stronger preference gets his way. If Y's preference for e over not-e is greater than X's preference for not-e over e, then e will occur; whereas not-e will occur if the converse holds. Similarly, if Y's preference for not-e over e exceeds S's for e over not-e, then not-e will occur; and e will occur if the converse holds. Here we may conclude that X's power and Y's power w.r.t. E are equal.

Since our primary interest is in comparisons of power, let us restrict our attention to northwest and southeast quadrants. Graphs 2 and 3 depict power-relationships where it is natural to say that Y has more power than X. Graph 2 says that Y gets his way no matter how small his preference and no matter how large X's opposing preference. In other words, Graph 2 says that Y is able to ensure either e or not-e at no cost whatever. (This would be true if the courses of action that would ensure e and not-e, respectively, had equal intrinsic attractiveness for Y and were each more attractive, apart from their effect on issue E, than any alternative course of action.) Graph 3 again implies, I believe, that Y has more power than X, but here his excess of power is not as great. According to this graph, there is a wide range of cases in which Y would get his preferred outcome with a smaller degree of preference than X, but there are no cases in which X would get his preferred outcome with a smaller degree of preference than Y. Other cases might also be mentioned in which one agent fairly clearly has more power than another. But I shall not pursue this further in this paper. There will be many cases, of course, in which overall comparisons of power are difficult to make. One such case is that of Matrix 1, which is depicted in graphical form in Graph 4. Here the pattern of power

is complicated, and it is not perfectly evident what conclusion should be drawn. No doubt the judgment we would intuitively make partly depends on the order of magnitude that the numbers in question convey. In any case, we should not expect our theory to yield precise judgments of power comparisons in all cases. Quite obviously, our pre-analytic judgments about power are not so precise to begin with, and we should heed Aristotle's advice not to expect (or impose) more precision than the subject-matter allows. What our theory does provide, however, is a framework for expressing in a systematic way the crucial factors that enter into an appraisal of the relative degree of power of two agents w.r.t. an issue.

Before leaving the topic of power and cost, I wish to mention a difficulty confronting our theory that I do not know how to handle. Our use of the notion of opportunity cost has one disadvantage. By measuring cost, and hence power, in terms of the best alternative course of action open to the agent, one makes the degree of cost depend upon actions open to the agent that may have no connection at all with issue E. Suppose, for example, that you can obtain outcome e with an expenditure of $ 100. I come along and offer to sell you for $ 100 an extremely valuable painting that you would love to own. According to our theory, it would follow that I have reduced your power w.r.t. issue E. If we assume, at any rate, that you only have $ 100 to spend (or less than $ 200), buying the painting and spending $ 100 to get outcome e will be mutually exclusive courses of action. Hence, when I present you with this new opportunity, I thereby increase the cost to you of obtaining outcome e. This consequence of our theory is somewhat counterintuitive (though not entirely counterintuitive, I think), but I do not know how to deal with it without abandoning our theory completely. It seems clear, however, that the notion of cost must play a central role in any adequate theory of power. The approach we have sketched here seems very promising, and it should not lightly be abandoned.

V. OVERALL POWER AND POWER "OVER" PEOPLE

The foregoing analysis has focused on the power that an agent or group of agents have w.r.t. a single issue. But the concern of most discussions of power in the political arena is not the power of individuals or groups vis-à-vis selected *single* issues, but rather the *overall* power of individuals or groups. We do not merely want to say that Melvin Laird or Nelson

Rockefeller has more power than I do w.r.t. this or that particular issue;
we want to say that each of these men has more power than I do *on the*
whole. On what basis, however, are such judgments to be made? Is it that
Rockefeller has power (or a higher degree of power) w.r.t. a larger *number*
of issues than I do? How are we supposed to *count* the number of issues
w.r.t. which each of us has power? Issues, it would seem, can be divided
up in any way one chooses, and hence both Rockefeller and I could be
said to have power w.r.t. indefinitely many issues. A better answer to our
question would appeal not to the number of issues over which we have
power, but to their *importance*. I may have power w.r.t. what I eat for
dinner tonight, what grades my students get, what sentences are written
on my blackboard, etc. But Rockefeller has power w.r.t. 'really important'
issues. But how is the 'importance' of an issue to be measured? One ele-
ment that should go into the determination of 'importance' (for present
purposes) is the number of people that would be affected by the issue. The
fact that I have more power than Rockefeller does w.r.t. what will be writ-
ten on my blackboard doesn't count for much in assessing my overall
power as compared with Rockefeller's. On the other hand, the fact that
Rockefeller has more power than I do w.r.t. the issues of taxation and
expenditure of the State of New York is a very important reason for say-
ing that he has more overall power. One thing that helps account for this
is that taxation and expenditure policies of the State of New York affect
a very large number of people.

But what do we mean when we say that an issue '*affects*' a person? And
can one issue 'affect' a man *more* or *less* than another issue? The answer to
the first question, I think, is suggested by looking at the second. I am in-
clined to say that one issue affects a person more than another issue if the
outcomes of the first issue would make a larger difference to his *welfare* than
the outcomes of the second. An issue whose outcomes would make no
difference whatever to his welfare does not affect him at all. In determining
a man's overall power, therefore, we must look not only at the number of
persons that would be affected by the issues w.r.t. which he has power,
but also at *how much* difference in welfare the outcomes of the issue would
cause. This principle accords nicely with our intuitions about the 'impor-
tant' issues of the day. The issue of war or peace, the issue of inflation and
unemployment, the issue of pollution and measures for combating it, the
issue of poverty and racism, are some of the major issues of contemporary

American life. They all have in common the two features we have mentioned: first, they all affect the welfare of a large number of people, and secondly, the outcomes of these issues make a very large difference in the welfare of these people. It is not surprising, therefore, that individuals or groups who have power w.r.t. these issues, and especially w.r.t. a large number of these issues, are regarded as having a great deal of overall power.

In determining a man's overall power, we shall not simply be interested in the *single* issues w.r.t. which he has power; we shall also be extremely interested in the question of his power w.r.t. *conjunctions* of issues. Suppose there are two independent issues, E and E^*, each partitioned into three outcomes. Issue E has outcomes e_1, e_2 and e_3, and issue E^* has outcomes e_1^*, e_2^* and e_3^*. Further suppose that Smith is in a position to ensure any of the outcomes of E through a payment of \$ 100 and is in a position to ensure any of the outcomes of E^* through a payment of \$ 100. Then if Smith has exactly \$ 100, he has (individual) power w.r.t. *each* of these issues. Now suppose, however, that we consider the *conjunctive* issue E & E^*. This issue has nine possible outcomes: e_1 & e_1^*, e_1 & e_2^*, ..., e_3 & e_2^*, e_3 & e_3^*. It does not follow from our original supposition that Smith has (individual) power w.r.t. this conjunctive issue. For if he only has \$ 100, he will not be in a position to ensure any of the nine indicated outcomes. If we suppose that Smith has \$ 200, however, it will follow that he not only has power w.r.t. each of the outcomes, E and E^*, but that he also has power w.r.t. the conjunctive issue, E & E^*. Obviously, in determining the economic and political power of men and groups, attention to conjunctions of issues is of crucial importance. If a man has power w.r.t. a conjunction of many issues, each of which significantly affects the welfare of many people, then he is a very powerful man.

In the foregoing paragraphs the notion of welfare has been taken for granted, and we shall continue to do this in the remainder of the paper. For illustrative purposes, let us associate a man's welfare with the satisfaction or non-satisfaction of his desires. (It is not clear that welfare should be identified with satisfaction of desires – rather than, say, with the satisfaction of needs or interests – but this is close enough for the moment.) Suppose that the degree of desire (or utility) assigned by Jones to the possible outcomes of E are as follows: $(e_1)+100$, $(e_2)+70$, $(e_3)+20$, $(e_4)-40$, and $(e_5)-100$. Further suppose that Jones has no power at all

w.r.t. issue E, but that Smith has enough power to ensure any one of the outcomes e_2, e_3, and e_4 (though not e_1 or e_5). We can then say that Smith has the power to make a fairly substantial difference in Jones' welfare — in particular, a difference of 110 utils (the difference between the utility of the best outcome for Jones that Smith can ensure and the utility of the worst outcome that Smith can ensure). Now if Smith had the power to ensure *any* of the five outcomes of E, and not just the middle three, then Smith's power over Jones would be even greater. Of course, our assessment of Smith's power over Jones is partly a function of the *degree* of power Smith has w.r.t. these outcomes. If it would be quite costly for Smith to secure any of these outcomes, then his power over Jones is less than it would be if the cost to Smith were small. Thus, in determining an agent's overall power, there are several factors to be considered: (1) the issues w.r.t. which the agent has *some* power, (2) the conjunctions of these issues w.r.t. which he has power, (3) the degree of power (in terms of cost) he has w.r.t. these issues, (4) the number of people affected by these issues, and (5) the amounts of the differences in these people's welfare that depend on the outcomes of these issues.

During most of our discussion we have concentrated on the notion of having power 'with respect to' an 'issue'. The term 'power', however, frequently occurs in sentences of the form 'X has power over Y', where 'Y' designates a person or group of people. We still owe an account, therefore, of having power *over* a *person*. Our remarks in the previous paragraph suggest such an account. Smith has power *over Jones*, we might say, if and only if Smith has power w.r.t. issues that affect Jones — i.e., that make a difference to Jones' welfare.

In addition to the support already given to this approach, the welfare accounts fits neatly with the relation between power and dependence that is often cited by sociologists.[26] If agent X controls certain objects or events which Y desires, or on which he sets a value, then Y is 'dependent' on X for the satisfaction of his desires or interests. Intuitively, this looks like a case in which X has power 'over' Y. This intuition is captured by the welfare approach we have sketched; for if Y is dependent on X for the satisfaction of his desires or interests, then X has power w.r.t. issues that affect the welfare of Y.

The suggested account of having power *over* people is in accord with a number of treatments of power in the literature. But many other writers

take a somewhat different approach. Writers such as R. H. Tawney, Robert Dahl, J. C. Harsanyi, and Herbert Simon conceive of power over people as the ability to affect the *behavior* of those people.[27] Tawney, for example, defines 'power' as "the capacity of an individual, or group of individuals, to modify the conduct of other individuals or groups in the manner which he desires..." (*op. cit.*, p. 230). Similarly, Robert Dahl says: "*A* has power over *B* to the extent that he can get *B* to do something that *B* would not otherwise do" (*op. cit.*, pp. 202–3).

Now as a complementary approach to the welfare account, the behavior approach is certainly welcome. But if it is intended to replace the welfare approach, I think it is unsatisfactory. Considered as a replacement, the behavior approach has three problems. First, it is difficult to see how it can throw much light on the determination of the *degrees* of power that person X has over person Y. The main possibility here would be to concentrate on the *number* of acts of Y w.r.t. which X has power.[28] But principles for counting acts or items of behavior are very controversial. On the theory I favor, most cases of changing the behavior of an individual would involve changes in *indefinitely many* acts. Moreover, on any approach to act-individuation one should be able to make temporal divisions any way one pleases, so it will be impossible to decide when one, two or three acts have been affected. The welfare approach, by contrast, holds greater promise for the measurement of degrees of power.

A second problem for the behavior approach is that behavior control is not clearly a sufficient condition for the exercise of power. Suppose that you are about to sit in a chair. I politely ask you to sit in the next chair instead, and since it is equally comfortable, you oblige. Here I have affected your behavior: I have induced you to do something you would not otherwise have done. But have I exercised *power over* you? This is doubtful. The reason it is doubtful, I think, can be explained by the welfare approach. First of all, there is no difference for your welfare between sitting in the one chair or sitting in the other. Secondly, since my method of inducing the change in behavior did not depend on threats or deprivation of opportunities, there is nothing in the case to suggest that I have any means of (importantly) affecting your welfare. A similar case is that of a sky-writer, whose tracks cause many people to perform acts they would not otherwise perform, viz., look up at the sky. There is little temptation to say that the sky-writer has exercised power over these

people. Although he causes a change in behavior, this change is unac-
companied by any (or very much) change in welfare.

A third problem for the behavior approach is the existence of cases
where behavior control fails but where there is inclination to talk of the
exercise of power. Suppose I threaten to beat you up unless you do act A.
You refuse to comply and so I beat you up. Here I have failed to control
your behavior, but isn't my beating you up itself an exercise of power?[29]
I am inclined to think that it is. My having power w.r.t. your getting
beaten up is a form of power over you, even when unaccompanied by
power w.r.t. your doing A. Again, this case is easily accounted for by the
welfare approach, since my having power w.r.t. the issue of your getting
beaten up is clearly having power w.r.t. an issue that importantly affects
your welfare.

It goes without saying, of course, that power w.r.t. behavior and power
w.r.t. welfare are in general closely interdependent. (This point is stressed
by many authors, including Harsanyi and Thibaut and Kelley to mention
just a few.) If X has power w.r.t. issues that affect Y's welfare, and if he
knows how they would affect Y's welfare, then he is in a position to make
threats or offers to Y in order to affect his behavior.[30] Moreover, certain
ways of affecting behavior characteristically result in behavior that is
less valuable to the agent affected. If X deters Y by threat from performing
an act, or if X prevents Y from performing an act he otherwise would have
performed, then there is typically a loss in utility for Y. Even though Y
averts the threatened sanction or penalty by complying with X's desires,
the substitute course of action is normally less valuable for Y than the
act he would have performed had X not intervened. For these reasons,
then, there is a close connection between power w.r.t. the *behavior* of Y
and power w.r.t. issues affecting the *welfare* of Y.

Nevertheless, it must be admitted that there are cases of power 'over'
people that involve only power w.r.t. their behavior, not power w.r.t.
their welfare. The most obvious cases of this sort are cases of what
Max Weber calls 'imperative control' or 'imperative coordination' (Herr-
schaft). When the commands of an authority in a corporate group are
accepted by the group as 'legitimate', there is a tendency to obey the
commands without the need for coercion (or reward). In many such cases,
moreover, compliance with these commands have no effect (or very little
effect) on the welfare of the compliers. For very often the behavior called

for by these imperatives is more-or-less routine behavior of functionaries which is quite inconsequential as far as they themselves are concerned. Needless to say, the possession of imperative control is normally an extremely valuable resource to have in controlling the welfare of people. That is, it normally *enables* its possessor to have significant degrees of power w.r.t. issues that affect people's welfare. The point, however, is that we can abstract from this power w.r.t. welfare and still be left with something important that would ordinarily be called power '*over*' people. Thus, it is important to retain the notion of behavior control as a distinct category of power 'over' people, despite its large overlap with the category of welfare control.

A third category of power 'over' people is also called for, though again the overlap with the other categories will be substantial. The third category is power to *persuade*, or more generally, power w.r.t. various psychological states of others, especially *desires, beliefs,* and *attitudes.* To the extent that a 'charismatic' leader has power w.r.t. the attitudes of his followers, and to the extent that the controller of communication media has power w.r.t. the desires and beliefs of the public, such men have power 'over' others. To be sure, power of this sort is closely connected with power w.r.t. behavior and power w.r.t. welfare. The difference between true and false belief, for example, will frequently make a difference to one's ability or inability to satisfy one's interests. Nevertheless, there are many cases in which the connection between desires and beliefs, on the one hand, and behavior and welfare, on the other, is at best very indirect and complicated. Many beliefs are never acted upon, and changes in desire are hard to relate to changes in welfare. For these reasons, then, it is worth including the third category as a distinct category of power 'over' others.

A further reason for emphasizing the third category must also be mentioned. Our analysis of power w.r.t. an issue, it should be recalled, makes an essential appeal both to hypothetical desires of an agent and to his actual desires (which determine the *cost* of an activity). It is obvious, therefore, that X's having power w.r.t. Y's desires can have an important bearing on Y's power w.r.t. an issue. In fact, this point raises an important question about our analysis of power w.r.t. issues. According to our analysis, we are interested in whether an outcome of an issue would occur *if* agent Y wanted that outcome. But suppose agent X is in a position to ensure that Y will *not* want this outcome. What does this imply about

Y's power w.r.t. that issue? This question must be taken up in a more fully developed theory of power, but this is not the occasion for such a development.[31]

The University of Michigan

NOTES

[1] A few of the most prominent works in the field are the following: Robert A. Dahl, 'The Concept of Power', *Behavioral Science* 2 (1957); J. R. P. French, Jr., 'A Formal Theory of Social Power', *Psychological Review* 63 (1956); H. Goldhamer and E. Shils, 'Types of Power and Status', *American Journal of Sociology* 45 (1939); J. C. Harsanyi, 'Measurement of Social Power, Opportunity Costs, and the Theory of Two-Person Bargaining Games', *Behavioral Science* 7 (1962); H. D. Lasswell and A. Kaplan, *Power and Society*, Yale University Press, New Haven (1950); J. G. March, 'An Introduction to the Theory and Measurement of Influence', *American Political Science Review*, 49 (1955); Bertrand Russell, *Power: A New Social Analysis*, New York: W. W. Norton (1938); L. S. Shapley and M. Shubik, 'A Method for Evaluating the Distribution of Power in a Committee System', *American Political Science Review* 48 (1954). For a critical review of relevant literature and some reservations about the utility of the power concept for empirical research, see J. G. March, 'The Power of Power', in D. Easton (ed.), *Varieties of Political Theory*, Prentice-Hall, Inc., Englewood Cliffs, N.J. (1966).

[2] See Robert Dahl, 'The Concept of Power', *Behavioral Science* 2 (1957), pp. 212–3. Actually, Dahl's chameleon just decides to *vote* with the predicted outcome, but I assume that his *desires* change to favor it as well.

[3] For further details on the notion of a basic act-type, see my book *A Theory of Human Action*, Prentice-Hall (1970), Chapter III, Section 4 (and also Chapter VII, Section 1). In this paper I am ignoring certain complications, such as the existence of 'standard conditions'. I shall also assume, for the most part, that all persons have the same basic act repertoire, though in one of my examples (in Section III) I shall drop this assumption.

[4] In ordinary language we frequently use the term 'power' without having a particular dated issue in mind. If John has a rain-making machine, for example, we might say that he has power over 'the weather', without specifying the time in question. This loose talk about power can be understood in terms of power w.r.t. dated issues. To say that John has power w.r.t. the weather (in general), is to say that there are dated issues involving the weather w.r.t. which John has power. Having a rain-making machine gives John power over the weather 'in general' because there are particular times at which it would not rain in the normal course of events, and w.r.t. rain *at these times* John has power. To say that John has power over the weather, in this sense, does not imply that John has power w.r.t. *every* dated issue involving the weather. If it is going to rain at time *t* no matter what John does (and if, in addition, he cannot control how hard it rains), then John has no power w.r.t. the issue of whether or not it rains *at t*. Assuming that he has no rain-prevention equipment, at any rate, he is impotent w.r.t. the issue of whether or not it rains *at t*.

[5] Actually, (I') needs still further refinement. First, it is not enough to require that *S* would perform the relevant acts; he must perform these acts *in order to* achieve the

given outcome (or at least *in the belief* that his performance of them would not prevent the occurrence of that outcome). Secondly, the analysis still remains open to counter-examples of the sort that beset (I). Suppose that S is not paralyzed and he falsely believes that whenever he snaps his fingers this causes rain clouds to disperse. It is true of S, then, that if he wanted non-rain he would snap his fingers (moreover, he would snap them in order to obtain non-rain). Further suppose, as in the earlier example, that S^* has the ability to detect S's desires by ESP, and he has resolved to make the weather conform with S's desire. Now although S's snapping his fingers would not be causally relevant to the occurrence of non-rain, there is a sequence of basic acts such that if S wanted non-rain he would perform these acts, and if he performed these acts non-rain would occur. Thus, S satisfies (I') w.r.t. the issue of rain or non-rain. But it is doubtful that we would credit S with power w.r.t. this issue.

What goes wrong here is that S is mistaken about the contribution that his finger-snapping would make to the occurrence of non-rain. He erroneously believes that this act would directly cause the dispersal of the clouds, and he is completely ignorant of the role of S^*. To circumvent such problems we need a provision requiring that the agent's *conception* of the relevance of his action to the outcome must correspond to the relevance his action *would* indeed have. I do not know how to formulate such a provision in adequate detail (similar problems are encountered in the analysis of intentional action and in the analysis of knowing), but I shall suggest a vague formulation for the present.

Each of the conjuncts of (I') should be amended as follows:
"There is a sequence of basic act-types such that   .

(a)   if S wanted e [not-e], then, at appropriate times between t_1 and t_n, he would perform these acts, either (i) in order to achieve outcome e [not-e], or (ii) in the belief that they would not preclude the occurrence of e [not-e], and

(b)   if S performed these acts at these appropriate times, then they would contribute to the occurrence of e [not-e] in the way S would expect them to contribute, and e [not-e] would occur (at t_n)."

(Notice that my reformulation of clause (b) does not imply that S's action would *cause* the outcome. This is because I want to allow, for one of the outcomes, that it would occur in the 'normal' course of events. This is all right as long as S *realizes*, for the appropriate outcome, that his action would merely 'allow' this outcome to occur, rather than *cause* it to occur.)

The indicated reformulation should be regarded as my 'official' analysis of individual power. I omit it from the text because I want to use (I') as a model for later analyses; to use the more refined analysis, I fear, would obscure the more important structural features of the theory.

[6] Cf. W. V. Quine, *Word and Object*, The Technology Press, Cambridge, Mass. (1960), pp. 222–6.

[7] Perhaps we should draw a distinction here between an 'epistemic' sense of 'power' and a 'non-epistemic' sense of 'power'. The epistemic sence is captured by (I') as it stands, whereas the non-epistemic sense can be obtained from (I') by deleting clauses (1) (a) and (2) (a). This distinction would parallel the distinction between the epistemic and non-epistemic senses of 'ability' that I draw in *A Theory of Human Action*, p. 203. For our purposes, however, the epistemic sense of 'power' seems much more important than the non-epistemic sense.

[8] Robert Stalnaker, 'A Theory of Conditionals', in Nicholas Rescher (ed.), *Studies in Logical Theory* (American Philosophical Quarterly Monograph Series), Basil Blackwell, Oxford (1968).

[8a] The difficulties here are well known. Cf. Nelson Goodman, *Fact, Fiction and Forecast*, Harvard University Press, Cambridge, Mass. (1955).

[9] It is necessary to rule out partitions with 'gimmicky' outcomes. For example, we do not want to partition the weather issue into the three outcomes: (a) it does not rain, (b) it rains with my right hand raised, (c) it rains with my right hand not raised. Unfortunately, I do not know how to formulate a general condition to rule out such a 'gimmicky' partition.

[10] Actually, a further condition – condition (d) – must be added. To see the necessity for this, suppose that a man will be chosen from a list of ten candidates, and let E be the issue of who will be chosen. Suppose we consider two partitions of the issue, P and P'. P is the partition whose outcomes are (i) a Democrat is chosen and (ii) a Republican is chosen. P' is the partition whose outcomes are (i') an Eastsider is chosen and (ii') a Westsider is chosen. (Both of these pairs are exhaustive and mutually exclusive.) Further suppose that in W_1 S could ensure either of the outcomes of P but not of P', whereas in W_2 S could ensure either of the outcomes of P' but not of P. Then it follows from (a), (b) and (c) that S has more power w.r.t. E in W_1 than in W_2 *and also* that he has more power w.r.t. E in W_2 than in W_1. To avoid this difficulty, the following condition, may be added:

(d)   There is no other partition P' of E, whose outcomes constitute set E^{**}, for which there are subsets E''' and E'''' such that E'''' contains at least two members, E''' is a proper subset of E'''', in W_1 E'''' is the largest subset of E^{**} any of whose members S can ensure, and in W_2 E''' is the largest subset of E^{**} any of whose members S can ensure.

[11] *New York Magazine*, Vol. IV, No. 1 (January 4, 1971), p. 25.

[12] *The Theory of Social and Economic Organization*, Oxford University Press, New York, (1947), p. 152. Italics are mine.

[13] *The Power Elite*, Oxford University Press, New York (1959), p. 9. Italics are mine.

[14] Cf. Thomas Schelling, *The Strategy of Conflict*, Harvard University Press, Cambridge, Mass. (1960), and David K. Lewis, *Convention*, Harvard University Press, Cambridge, Mass. (1969), Chapter 1.

[15] Although Smith does not have individual power w.r.t. the issue of whether or not Brown is alive at noon, he does have individual power w.r.t. the (distinct) issue of whether or not *he* (Smith) *kills* Brown. Thus, even apart from our distributive principle, we have a way of accounting for the intuition that Smith has *some* power here.

[16] In the game whose matrix is shown below, Column and Row make one move each, simultaneously, with no information about the move of his opponent. We can say here that for any move by Column there is a move available to Row which would 'win' for Row; but similarly, for any move by Row there is a move available to Column which would 'win' for Column.

	c_1	c_2
r_1	$+100$	-100
r_2	-100	$+100$

[17] In typical political situations it will be necessary to partition an issue into numerous

possible outcomes, e.g., the different possible compromises that might be reached on a piece of legislation, or the different wage increases that might be given to a union. Suppose there is a partition of E into 10 outcomes, and suppose that if S_1 and S_2 had diametrically opposite rank orderings then the outcome which is fourth on S_1's list (and seventh on S_2's list) would occur. Does it follow that S_1 has more power than S_2 w.r.t. E? This is not evident (even neglecting other partitions). Further complications arise if S_1 and S_2 have different but not diametrically opposite rank orderings.

[18] Reflection on these kinds of cases suggests that statements expressing comparisons of power might be parsed, not in terms of a two-place relation 'x has more power than y', but rather in terms of a three- (or more) place relation such as 'x has more power than y given z as a resource'. This sort of treatment would make explicit which person(s) are being treated as resources. It would also have the virtue of forestalling certain problems of transitivity that might otherwise arise.

[19] I assume that the machine can only be operated by S^*; thus, S can only affect the weather *through* S^*, not by operating the machine himself.

[20] 'Measurement of Social Power, Opportunity Costs, and the Theory of Two-Person Bargaining Games', *Behavioral Science* 7 (1962), pp. 67–80.

[21] This point raises difficult problems. If we counterfactualize S's desires vis-à-vis E, it may seem plausible to counterfactualize a variety of other desires of his as well. It would not be plausible to consider the (counterfactual) hypothesis that Nixon wants a Democrat to be elected in 1972 without supposing other significant changes in his set of desires. The difficulty here is to decide which changes to make. There is no unique set of changes one is forced to make; numerous counterfactualizations are equally admissible. This is undoubtedly one of the reasons why judgments about power in ordinary life are so ambiguous and difficult to agree about: it is possible to counterfactualize in many different ways, yet these different alternatives are not normally made explicit. For the sake of simplicity, we shall assume in the remainder of the text that *no* additional changes are made in the agent's desires beyond the changes in his desire vis-à-vis E.

[22] Strictly speaking, it is not the *actual* costs to be undergone by a person that are relevant to his power; rather, it is his *expected* costs. To avoid any further complications, however, let us assume that both S_1 and S_2 have perfect information about the consequences of their activities, so that expected and actual costs are the same.

[23] A question here arises whether to include the (expected) effects of e and of not-e in determining the cost of an activity or whether to include these under the utility assigned to e and not-e themselves (hence omitting them from the determination of cost). This is a tricky question, since the degree to which a given outcome of E is desired commonly depends upon expected consequences of that outcome. What I propose is the following. If C is expected to result from outcome e, then if C would constitute one of the agent's *reasons* for wanting e, then it is included under the value of e (*not* in the category of *cost*). If C would be an *unwelcome* consequence of e (e.g., going to jail), then it is included in the category of cost.

[24] For convenience we assume, in this case, that the utility accruing to Row from this course of action is always the same, no matter which sequence of acts is chosen by Column.

[25] The introduction of many-outcome partitions also provides a tool for distinguishing degrees of power; but the new approach, making use of degrees of desire, is a more powerful one, I believe.

[26] Cf. J. W. Thibaut and H. H. Kelley, *The Social Psychology of Groups*, Wiley and