

TWO-PERSON COOPERATIVE GAMES¹

BY JOHN NASH

In this paper, the author extends his previous treatment of "The Bargaining Problem" to a wider class of situations in which threats can play a role. A new approach is introduced involving the elaboration of the threat concept.

INTRODUCTION

THE THEORY presented here was developed to treat economic (or other) situations involving two individuals whose interests are neither completely opposed nor completely coincident. The word cooperative is used because the two individuals are supposed to be able to discuss the situation and agree on a rational joint plan of action, an agreement that should be assumed to be enforceable.

It is conventional to call these situations "games" when they are being studied from an abstract mathematical viewpoint. Here the original situation is reduced to a mathematical description, or model. In the abstract "game" formulation only the minimum quantity of information necessary for the solution is retained. What the actual alternative courses of action are among which the individuals must choose is not regarded as essential information. These alternatives are treated as abstract objects without special qualities and are called "strategies." Only the attitudes (like or dislike) of the two individuals towards the ultimate results of the use of the various possible opposing pairs of strategies are considered; but this information must be well utilized and must be expressed quantitatively.

The theory of von Neumann and Morgenstern applies to some of the games considered here. Their assumption that it is possible for the players to make "side-payments" in a commodity for which each individual (player) has a linear utility narrows the range of their theory's applicability. In this paper there is no assumption about side-payments. If the situation permits side-payments then this simply affects the set of possible final outcomes of the game; side-payments are treated just like any other activity that may take place in the actual playing of the game—no special consideration is necessary. The von Neumann and Morgenstern approach also differs by giving a much less determinate solution. Their approach leaves the final situation only determined up to a side-payment. The side-payment is generally not determined but is restricted to lie in a certain range.

An earlier paper by the author [3] treated a class of games which are in one sense the diametrical opposites of the cooperative games. A game is

¹ This paper was written with the support of The RAND Corporation. It appeared in an earlier form as RAND P-172, August 9, 1950.

non-cooperative if it is impossible for the players to communicate or collaborate in any way. The non-cooperative theory applies without change to any number of players, but the cooperative case, which is analyzed in this paper, has only been worked out for two players.

We give two independent derivations of our solution of the two-person cooperative game. In the first, the cooperative game is reduced to a non-cooperative game. To do this, one makes the players' steps of negotiation in the cooperative game become moves in the non-cooperative model. Of course, one cannot represent all possible bargaining devices as moves in the non-cooperative game. The negotiation process must be formalized and restricted, but in such a way that each participant is still able to utilize all the essential strengths of his position.

The second approach is by the axiomatic method. One states as axioms several properties that it would seem natural for the solution to have and then one discovers that the axioms actually determine the solution uniquely. The two approaches to the problem, via the negotiation model or via the axioms, are complementary; each helps to justify and clarify the other.

THE FORMAL REPRESENTATION OF THE GAME

Each of the players (one and two) has a compact convex metrizable space S_i of mixed strategies s_i (those readers who are unacquainted with the mathematical technicalities will find that they can manage quite well by ignoring them). These mixed strategies represent the courses of action player i can take independently of the other player. They may involve deliberate decisions to randomize, to decide between alternative possibilities by using a randomizing process involving specified probabilities. This randomizing is an essential ingredient in the concept of a mixed strategy. By beginning with a space of mixed strategies instead of talking about a sequence of moves, etc., we presuppose a reduction of the strategic potentialities of each player to the normal form [4].

The possible joint courses of action by the players would form a similar space. But the only important thing is the set of those pairs (u_1, u_2) of utilities which can be realized by the players if they cooperate. We call this set B and it should be a compact convex set in the (u_1, u_2) plane.

For each pair (s_1, s_2) of strategies from S_1 and S_2 , there will be the utility to each player of a situation where these strategies are to be employed or carried out. These utilities (pay-offs in game theoretic usage) are denoted by $p_1(s_1, s_2)$ and $p_2(s_1, s_2)$. Each p_i is a linear function of s_1 and of s_2 , although it cannot be expected to depend linearly on the two varying simultaneously; in other words, p_i is a bilinear function of s_1 and s_2 . Basically this linearity is a consequence of the type of utility we assume for the players; it is thoroughly discussed in an early chapter of von Neumann and Morgenstern [4].

And of course each point in the (u_1, u_2) plane of the form $[p_1(s_1, s_2), p_2(s_1, s_2)]$ must be a point in B because every pair (s_1, s_2) of independent strategies corresponds to a joint policy (probably an inefficient one). This remark completes the formal, or mathematical, description of the game.

THE NEGOTIATION MODEL

To explain and justify the negotiation model used to obtain the solution we must say more about the general assumptions about the situation facing the two individuals, or, what it amounts to, about the conditions under which the game is to be played.

Each player is assumed fully informed on the structure of the game *and* on the utility function of his co-player (of course he also knows his own utility function). (This statement must not be construed as inconsistent with the indeterminacy of utility functions up to transformations of the form $u' = au + b$, $a > 0$.) These information assumptions should be noted, for they are not generally perfectly fulfilled in actual situations. The same goes for the further assumption we need that the players are intelligent, rational individuals.

A common device in negotiation is the threat. The threat concept is really basic in the theory developed here. It turns out that the solution of the game not only gives what should be the utility of the situation to each player, but also tells the players what threats they should use in negotiating.

If one considers the process of making a threat, one sees that its elements are as follows: A threatens B by convincing B that if B does not act in compliance with A 's demands, then A will follow a certain policy T . Supposing A and B to be rational beings, it is essential for the success of the threat that A be *compelled* to carry out his threat T if B fails to comply. Otherwise it will have little meaning. For, in general, to execute the threat will not be something A would want to do, just of itself.

The point of this discussion is that we must assume there is an adequate mechanism for forcing the players to stick to their threats and demands once made; and one to enforce the bargain, once agreed. Thus we need a sort of umpire, who will enforce contracts or commitments.

And in order that the description of the game be complete, we must suppose that the players have no prior commitments that might affect the game. We must be able to think of them as completely free agents.

THE FORMAL NEGOTIATION MODEL

Stage one: Each player (i) chooses a mixed strategy t_i which he will be forced to use if the two cannot come to an agreement, that is, if their demands are incompatible. This strategy t_i is player i 's threat.

Stage two: The players inform each other of their threats.

Stage three: In this stage the players act independently and without communication. The assumption of independent action is essential here, whereas no special assumptions of this type are needed in Stage one, as it turns out. In Stage three, each player decides upon his demand d_i , which is a point on his utility scale. The idea is that player i will not cooperate unless the mode of cooperation has at least the utility d_i to him.

Stage four: The pay-offs are now determined. If there is a point (u_1, u_2) in B such that $u_1 \geq d_1$ and $u_2 \geq d_2$, then the pay-off to each player i is d_i . That is, if the demands can be simultaneously satisfied, then each player gets what he demanded. Otherwise, the pay-off to player i is $p_i(t_1, t_2)$; i.e., the threats must be executed.

The choice of the pay-off function in the case of compatible demands may seem unreasonable, but it has its advantages. It cannot be accused of contributing a bias to the final solution and it gives the players a strong incentive to increase their demands as much as is possible without losing compatibility. But it can be embarrassingly accused of picking points that are not in the set B . Effectively, we have enlarged B to a set including all utility pairs dominated (weakly; $u'_1 \leq u_1, u'_2 \leq u_2$) by a pair in B .

What we have is actually a two move game. Stages two and four do not involve any decisions by the players. The second move choices are made with full information about what was done in the first move. Therefore, the game consisting of the second move alone may be considered separately (it is a game with a variable pay-off function determined by the choices made at the first move). The effect of the choice of threats on this game is to determine the pay-offs if the players do not cooperate.

Let N be the point $[p_1(t_1, t_2), p_2(t_1, t_2)]$ in B . This point N represents the effect of the use of the threats. Let u_{1N} and u_{2N} abbreviate the coordinates of N . If we introduce a function $g(d_1, d_2)$ which is +1 for compatible demands and 0 for incompatible demands, then we can represent the pay-offs as follows:

$$\begin{array}{ll} \text{to player one} & d_1g + u_{1N}(1 - g), \\ \text{to player two} & d_2g + u_{2N}(1 - g). \end{array}$$

The demand game defined by these pay-off functions will generally have an infinite number of inequivalent equilibrium points [3]. Every pair of demands which graphs as a point on the upper-right boundary of B and which is neither lower nor to the left of N will form an equilibrium point. Thus the equilibrium points do not lead us immediately to a solution of the game. But if we discriminate between them by studying

their relative stabilities we can escape from this troublesome non-uniqueness.

To do this we "smooth" the game to obtain a continuous pay-off function and then study the limiting behavior of the equilibrium points of the smoothed game as the amount of smoothing approaches zero.

A certain general class of natural smoothing methods will be considered here. This class is broader than one might at first think, for many other methods that superficially seem different are actually equivalent.

To smooth the game we approximate the discontinuous function g by a continuous function h , which has a value near to g 's value except at the points near the boundary of B , where g is discontinuous. The function $h(d_1, d_2)$ should be thought of as representing the probability of compatibility of the demands d_1 and d_2 . It can be thought of as representing uncertainties in the information structure of the game, the utility scales, etc. For convenience, let us assume that $h = 1$ on B and that h tapers off very rapidly towards zero as (d_1, d_2) moves away from B , without ever actually reaching zero. Another simplification can be had by assuming the utility functions properly transformed so that $u_{1N} = u_{2N} = 0$. Then we can write the pay-off functions for the smoothed game as $P_1 = d_1h$, $P_2 = d_2h$. For the original game h is replaced by g .

A pair of demands (d_1, d_2) viewed as a pair of pure strategies in the demand game, will be an equilibrium point if p_1 , which is d_1h , is maximized here for constant d_2 and if $p_2 = d_2h$ is maximized for constant d_1 . Now suppose (d_1, d_2) is a point where d_1d_2h is maximized over the whole region in which d_1 and d_2 are positive. Then d_1h and d_2h will be maximized for constant d_2 and d_1 , respectively, and (d_1, d_2) must be an equilibrium point.

If the function h decreases with increasing distance from B in a wavy or irregular way, there may be more equilibrium points and perhaps even more points where d_1d_2h is a maximum. But if h varies regularly there will be only one equilibrium point coinciding with a unique maximum of d_1d_2h . However, we do not need to appeal to a regular h to justify the solution.

Let P be any point where d_1d_2h or, what is the same thing, u_1u_2h is maximized as above described and let ρ be the maximum of u_1u_2 on the part of B lying in the region $u_1 \geq 0$, $u_2 \geq 0$. The value of u_1u_2 at P must be at least ρ , since $0 \leq h \leq 1$ and since $h = 1$ on B . Figure 1 illustrates this situation. In it, Q is the point where u_1u_2 is maximized on B (in the first quadrant about N) and $\alpha\beta$ is the hyperbola $u_1u_2 = \rho$, which touches B at Q .

The important observation is that P must lie above $\alpha\beta$ but still be near enough to B for h to nearly equal 1. And as less and less smoothing

is used, h will decrease more and more rapidly on moving away from B ; hence any maximum point P of $u_1 u_2 h$ will have to be nearer and nearer to B . In the limit all such points must approach Q , the only contact point of B and the area above $\alpha\beta$. Thus Q is a necessary limit of equilibrium points, and Q is the only one.

We take Q for the solution of the demand game, characterized as the *only necessary limit of the equilibrium points of smoothed games*. The values of u_1 and u_2 at Q will be taken as the values of the demand game and as the optimal demands.




FIGURE 1

The discussion above implicitly assumed that B contained points where $u_1 > 0$, $u_2 > 0$ (after the normalization which made $u_1 = u_2 = 0$ at N). The other cases can be treated more simply without resource to a smoothing process. In these "degenerate cases" there is only one point of B which dominates the point N and is not itself dominated by some other point of B . (A point (u_1, u_2) is dominated by another point (u'_1, u'_2) if $u'_1 \geq u_1$ and $u'_2 \geq u_2$) [see Figure 3]. This gives us the natural solution in these cases.

One should note that the solution point Q of the demand game varies as a continuous function of the threat point N . Also there is a helpful geometrical characterization of the way Q depends on N . The solution point Q is the contact point with B of a hyperbola whose asymptotes are the vertical and horizontal lines through N . Let T be the tangent at Q to this hyperbola (see Figure 2).

If linear transformations are applied to the utility functions, N can be made the origin and Q the point $(1, 1)$. Now T will have slope -1 and the line NQ will have slope $+1$. The essential point is that slope $T = \text{minus slope } NQ$, because this is a property that is not destroyed by linear transformations of the utilities. T will be a support line for the set B (that is, a line such that all points of B are either on the lower left side of T or are on T itself; for a proof, see reference [2] where the same situation arises).

We can now state the criterion: if NQ has positive slope and a support line T for B passes through Q with a slope equal but opposite to NQ 's slope, then Q is the solution point for the threat point N . If NQ is horizontal/vertical and is itself a support line for B and if Q is




FIGURE 2

the rightmost/uppermost of the points common to B and NQ , then again Q is the solution point for N (see Figure 3), and one of these cases must hold if Q is N 's solution point; the criterion is a necessary and sufficient one.

Any support line of B with a contact point Q on the upper-right boundary of B determines a complementary line through Q with equal but opposite slope. All points on the line segment in which this complementary line intersects B are points which, as threat points, would have Q as corresponding solution point. The class of all these line segments is a ruling of B by line segments which intersect, if at all, only on the upper-right boundary of B . Given a threat point N , its solution point is the upper-right end of the segment passing through it (unless perhaps N is on more than one ruling and hence is on the upper-right boundary and is its own solution point).

We can now analyze the threat game, the game formed by the first move and with pay-off function determined by the solution of the demand game. This pay-off is determined by the location of N , specifically by the ruling on which N falls. A ruling that is higher (or farther left) is more favorable to player two (let us definitely think of u_2 as measured on the vertical axis of the utility plane) and less favorable to player one.

Now if one player's threat is held fixed, say player one's at t_1 , then the position of N is a function of the other player's threat, t_2 . The coordinates of N , $p_1(t_1, t_2)$ and $p_2(t_1, t_2)$ are linear functions of t_2 . Hence the transformation, t_2 goes into N , defined by this situation is a linear transformation of the space S_2 of player two's threats into B . That part of the image of S_2 that falls on the most favorable (for player two) ruling will contain the images of the threats that would be best as replies to player one's fixed particular threat t_1 . And this set of best replies




FIGURE 3

must be a convex compact subset of S_2 because of the linearity and continuity of the transformation of S_2 into B .

The continuity of N as a function of t_1 and t_2 and the continuity of Q as a function of t_1 insure that the pay-off function defined for the threat game by solving the demand game is a continuous function of the threats. And this suffices to make each player's set of best replies what is called an upper semi-continuous function of the threat being replied to. Now consider any pair of threats (t_1, t_2) . For each threat of the pair the other player has a set of best replies. Let $R(t_1, t_2)$ be the set of all pairs which contain one threat from each of the two sets of replies. R will be an upper semi-continuous function of (t_1, t_2) in the space of opposed pairs of threats, and $R(t_1, t_2)$ will always be a convex set in this space, $S_1 \times S_2$.

We are now ready to use the Kakutani fixed point theorem, as generalized by Karlin [1, pp. 159–160]. This theorem tells us that there is some pair (t_{10}, t_{20}) that is contained in its set $R(t_{10}, t_{20})$, which amounts to saying that each threat is a best reply to the other. Thus we have obtained an equilibrium point in the threat game. It is worth noting that this equilibrium point is formed by pure strategies in the threat game

(a mixed strategy here would involve randomization over several threats).

The pair (t_{10}, t_{20}) also has minimax and maximin properties. Since the final pay-off in the game is determined by the position of Q on the upper-right boundary of B , which is a negatively sloping curve, each player's pay-off is a monotone decreasing function of the others. So if player one sticks to t_{10} , player two cannot make one worse off than he does by using t_{20} without improving his own position, and he can't do this because (t_{10}, t_{20}) is an equilibrium point [3]. Thus t_{10} assures player one the equilibrium pay-off and t_{20} accomplishes the same for player two.

The threat game is now revealed to be very much like a zero-sum game, and one can readily see that if one player were to choose his threat first and inform the other, rather than their simultaneously choosing threats, this would not make any difference because there is a "saddle-point" in pure strategies. It is rather different with the demand game. The right to make the first demand would be quite valuable, so the simultaneity here is essential.

To summarize, we have now solved the negotiation model, found the values of the game to the two players, and shown that there are optimal threats and optimal demands (the optimal demands are the values).

THE AXIOMATIC APPROACH

Rather than solve the two-person cooperative game by analyzing the bargaining process, one can attack the problem axiomatically by stating general properties that "any reasonable solution" should possess. By specifying enough such properties one excludes all but one solution.

The axioms below lead to the same solution that the negotiation model gave us; yet the concepts of demand or threat do not appear in them. Their concern is solely with the relationship between the solution (interpreted here as the value) of the game and the basic spaces and functions which give the mathematical description of the game.

It is rather significant that this quite different approach yields the same solution. This indicates that the solution is appropriate for a wider variety of situations than those which satisfy the assumptions we made in the approach via the model.

The notation used below is the same as before, except for a few additions. A triad (S_1, S_2, B) stands for a game and $v_1(S_1, S_2, B)$ and $v_2(S_1, S_2, B)$ are its values to the two players. Of course the triadic representation, (S_1, S_2, B) , leaves implicit the payoff functions $p_1(s_1, s_2)$ and $p_2(s_1, s_2)$ which must be given to determine a game.

AXIOM I: For each game (S_1, S_2, B) there is a unique solution (v_1, v_2) which is a point in B .

AXIOM II: If (u_1, u_2) is in B and $u_1 \geq v_1$ and $u_2 \geq v_2$ then $(u_1, u_2) = (v_1, v_2)$; that is, the solution is not weakly dominated by any point in B except itself.

AXIOM III: Order preserving linear transformations of the utilities ($u'_1 = a_1u_1 + b_1$, $u'_2 = a_2u_2 + b_2$ with a_1 and a_2 positive) do not change the solution. It is understood that the numerical values will be changed by the direct action of the utility transformations, but the relative position of (v_1, v_2) in B should stay the same.

AXIOM IV: The solution does not depend on which player is called player one. In other words, it is a symmetrical function of the game.

AXIOM V: If a game is changed by restricting the set B of attainable pairs of utilities and the new set B' still contains the solution point of the original game, then this point will also be the solution point of the new game. Of course the new set B' must still contain all points of the form $[p_1(s_1, s_2), p_2(s_1, s_2)]$, where s_1 and s_2 range over S_1 and S_2 , to make (S_1, S_2, B') a legitimate game.

AXIOM VI: A restriction of the set of strategies available to a player cannot increase the value to him of the game. Symbolically, if S'_1 is contained in S_1 , then $v_1(S'_1, S_2, B) \leq v_1(S_1, S_2, B)$.

AXIOM VII: There is some way of restricting both players to single strategies without increasing the value to player one of the game. In symbols, there exist s_1 and s_2 such that $v_1(s_1, s_2, B) \leq v_1(S_1, S_2, B)$. Similarly, there is a way to do the same for player two.

There is little need to comment on Axiom I; it is just a statement on the type of solution desired. Axiom II expresses the idea that the players should succeed in cooperating with optimal efficiency. The principle of noncomparability of utilities is expressed in Axiom III. Each player's utility function is regarded as determined only up to order preserving linear transformations. This indeterminacy is a natural consequence of the definition of utility [4, chapter 1, part 3]. To reject Axiom III is to assume that some additional factor besides each individual's relative preferences for alternatives is considered to make the utility functions more determinate and to assume that this factor is significant in determining the outcome of the game.

The symmetry axiom, Axiom IV, says that the only significant (in determining the value of the game) differences between the players are those which are included in the mathematical description of the game, which includes their different sets of strategies and utility functions. One may think of Axiom IV as requiring the players to be intelligent and rational beings. But we think it is a mistake to regard this as expressing "equal bargaining ability" of the players, in spite of a

statement to this effect in "The Bargaining Problem" [2]. With people who are sufficiently intelligent and rational there should not be any question of "bargaining ability," a term which suggests something like skill in duping the other fellow. The usual haggling process is based on imperfect information, the hagglers trying to propagandize each other into misconceptions of the utilities involved. Our assumption of complete information makes such an attempt meaningless.

It is probably harder to give a good plausibility argument for Axiom V than for any of the others. There is some discussion of it in "The Bargaining Problem" [2]. This axiom is equivalent to an axiom of "localization" of the dependence of the solution point on the shape of the set B . The location of the solution point on the upper-right boundary of B is determined only by the shape of any small segment of the boundary that extends to both sides of it. It does not depend on the rest of the boundary curve.

Thus there is no "action at a distance" in the influence of the shape of B on the location of the solution point. Thinking in terms of bargaining, it is as if a proposed deal is to compete with small modifications of itself and that ultimately the negotiation will be understood to be restricted to a narrow range of alternative deals and to be unconcerned with more remote alternatives.

The last two axioms are the only ones that are primarily concerned with the strategy spaces S_1 and S_2 , and the only ones that are really new. The other axioms are simply appropriate modifications of the axioms used in "The Bargaining Problem." Axiom VI says that a player's position in the game is not improved by restricting the class of threats available to him. This is surely reasonable.

The need for Axiom VII is not immediately obvious. Its effect is to remove the possibility that the value to a player of his space of threats should be dependent on collective or mutual reinforcement properties of the threats. The way Axiom VII is used in the demonstration of the adequacy of the axioms probably reveals its real content better than any heuristic discussion we might give here.

We can shortcut some of the arguments needed to show that the axioms accomplish their purpose and characterize the same solution we obtained with the model by appealing to the results of "The Bargaining Problem." We first consider games where each player has but one possible threat. Such a game is essentially a "bargaining problem," and for that sort of game our Axioms I, II, III, IV, and V are the same as the axioms of "The Bargaining Problem."

This determines the solution in the case where each player has but one strategy available. It must be the same solution obtained in "The Bargaining Problem," which was the same as the solution we got for

the demand game (which is played after each player has chosen a threat) in the preceding approach. This solution is characterized by the maximization of the product, $[v_1 - p_1(t_1, t_2)] [v_2 - p_2(t_1, t_2)]$, of the differences between the values of the game and the utilities of the situation where the players do not cooperate.

However, we are obliged to remark that the situation to be treated here is more general than that in "The Bargaining Problem" because it was assumed in that paper that there was some way for the players to cooperate with mutual benefit. Here it may be the case that only one, or neither, of the players can actually gain by cooperation. To show that the axioms handle this case seems to require a more complicated argument using Axioms VI and VII. But this is a minor point and we shall not include that argument, which is long out of proportion to its significance.

The primary function of Axioms VI and VII is to enable us to reduce the problem of games where each player may have a non-trivial space of strategies (threats) to the case we have just dealt with, where each has but one possible threat. Suppose player one is restricted to a strategy t_{10} which would be an optimal threat in the threat game discussed before in the non-axiomatic approach. Then from Axiom VI, we have

$$v_1(t_{10}, S_2, B) \leq v_1(S_1, S_2, B).$$

Now we apply Axiom VII to restrict S_2 to a single strategy (S_1 is already restricted) without increasing the value of the game to player one. Let t_2^* stand for the single strategy that S_2 is restricted to, then

$$v_1(t_{10}, t_2^*, B) \leq v_1(t_{10}, S_2, B).$$

Now we know that the value of a game where each player has but one threat is the same value obtained in the first part of this paper. Hence we know that against the threat t_{10} there is no better threat for player two, and no threat more unfavorable for player one, than t_{20} (i.e., an optimal threat for player two). So we may write

$$v_1(t_{10}, t_{20}, B) \leq v_1(t_{10}, t_2^*, B).$$

Combining the three inequalities we have

$$v_1(t_{10}, t_{20}, B) \leq v_1(S_1, S_2, B).$$

Similarly, we have

$$v_2(t_{10}, t_{20}, B) \leq v_2(S_1, S_2, B).$$

And now we observe, by Axiom II, that the last two inequalities may be replaced by equalities, because $v_1(t_{10}, t_{20}, B)$ and $v_2(t_{10}, t_{20}, B)$ are the

coordinates of a point on the upper-right boundary of B. Thus the axiomatic approach gives the same values as the other approach.

Massachusetts Institute of Technology

REFERENCES

- [1] KUHN, H. W. AND A. W. TUCKER, eds., *Contributions to the Theory of Games* (*Annals of Mathematics Study No. 24*), Princeton: Princeton University Press, 1950, 201 pp.
- [2] NASH, JOHN, "The Bargaining Problem," *ECONOMETRICA*, Vol. 18, April, 1950, pp. 155-162.
- [3] NASH, JOHN, "Non-Cooperative Games," *Annals of Mathematics*, Vol. 54, September, 1951, pp. 286-295.
- [4] VON NEUMANN, J. AND O. MORGENTERN, *Theory of Games and Economic Behavior*, 2nd edition, Princeton: Princeton University Press, 1947, 641 pp.