

Cambridge University Press

0521021332 - The Shapley Value: Essays in Honor of Lloyd S. Shapley

Edited by Alvin E. Roth

Excerpt

[More information](#)

CHAPTER 1

Introduction to the Shapley value

Alvin E. Roth

Among the obligations facing a community of scholars is to make accessible to a wider community the ideas it finds useful and important. A related obligation is to recognize lasting contributions to ideas and to honor their progenitors. In this volume we undertake to fill part of both obligations.

The papers in this volume review and continue research that has grown out of a remarkable 1953 paper by Lloyd Shapley. There he proposed that it might be possible to evaluate, in a numerical way, the “value” of playing a game. The particular function he derived for this purpose, which has come to be called the Shapley value, has been the focus of sustained interest among students of cooperative game theory ever since. In the intervening years, the Shapley value has been interpreted and reinterpreted. Its domain has been extended and made more specialized. The same value function has been (re)derived from apparently quite different assumptions. And whole families of related value functions have been found to arise from relaxing various of the assumptions.

The reason the Shapley value has been the focus of so much interest is that it represents a distinct approach to the problems of complex strategic interaction that game theory seeks to illuminate. To explain this, we need to recount some history of game theory. (Even when we are not speaking of the Shapley value, the history of game theory is inextricably connected with other aspects of Shapley’s work. To avoid too many diversions, we defer discussion of Shapley’s other work until the concluding section of this introduction.)

Although game-theoretic ideas can be traced earlier, much of the modern theory of games traces its origins to the monumental 1944 book by John von Neumann and Oskar Morgenstern, *Theory of games and economic behavior*. In seeking a way to analyze potentially very complex patterns of strategic behavior, their approach was to, in their phrase,

Cambridge University Press

0521021332 - The Shapley Value: Essays in Honor of Lloyd S. Shapley

Edited by Alvin E. Roth

Excerpt

[More information](#)

2 Alvin E. Roth

“divide the difficulties,” by finding simple models of the strategic environment itself.

Their first step was to find a way to summarize each alternative facing an individual decision maker by a single number. Their solution to this problem – expected utility theory – has left its own indelible mark on economic theory, quite independently of the impact the theory of games has had. Briefly, their contribution was to specify conditions on an individual’s preferences over possibly risky alternatives sufficient so that his choice behavior could be modeled as if, faced with a choice over any set of alternatives, he chose the one that maximized the expected value of some real-valued function, called his utility function. In this way, a complex probability distribution over a diverse set of alternatives could be summarized by a single number, equal to the expected utility of the lottery in question.

Having reduced the alternatives facing each individual to a numerical description, von Neumann and Morgenstern proceeded to consider (among other things) a class of games in which the opportunities available to each coalition of players could also be described by a single number. They considered cooperative games in *characteristic function form* (now sometimes also called “coalitional form”) defined by a finite set $N = \{1, \dots, n\}$ of players, and a real-valued “characteristic function” v , defined on all subsets of N (with $v(\emptyset) = 0$). The interpretation of v is that for any subset S of N the number $v(S)$ is the worth of the coalition, in terms of how much “utility” the members of S can divide among themselves in any way that sums to no more than $v(S)$ if they all agree. The only restriction on v that von Neumann and Morgenstern proposed was that it be superadditive; that is, if S and T are two disjoint subsets of N , then $v(S \cup T) \geq v(S) + v(T)$. This means that the worth of the coalition $S \cup T$ is equal to at least the worth of its parts acting separately.

The characteristic function model assumes the following things about the game being modeled. First, utility can be embodied in some medium of exchange – “utility money” – that is fully transferable among players, and such that an additional unit of transferable utility always adds a unit to any player’s utility function. (For example, if all players are risk neutral in money – that is, if their utility functions are all linear in money – then ordinary money can be the necessary medium of exchange in a game in which all outcomes can be evaluated in monetary terms and in which money is freely transferable.) Second, the possibilities available to a coalition of players can be assessed without reference to the players not included in the coalition. Third, a coalition can costlessly make binding

Cambridge University Press

0521021332 - The Shapley Value: Essays in Honor of Lloyd S. Shapley

Edited by Alvin E. Roth

Excerpt

[More information](#)**Introduction to the Shapley value**

3

agreements to distribute its worth in any way agreed to by all the members, so it is not necessary to model explicitly the actions that players must take to carry out these agreements. In recognition of the importance of the assumption that utility is transferable, these games are sometimes called transferable utility (TU) games.

Although these simplifying assumptions are obviously substantial, the characteristic function model has proved to be surprisingly useful as a simple model of strategic interaction. Consider, for example, the interaction between a potential seller and two potential buyers of some object that the seller (the current owner) values at ten dollars, the first buyer values at twenty dollars, and the second buyer values at thirty dollars. If the players can freely transfer money among themselves, and if they are risk neutral (although for many purposes this latter assumption is not really necessary), this situation can be modeled as the game $\Gamma_1 = (N, v)$ with players $N = \{1, 2, 3\}$ and v given by $v(1) = 10$, $v(2) = v(3) = v(23) = 0$, $v(12) = 20$, $v(13) = v(123) = 30$. This reflects the fact that only coalitions containing the seller, player 1, and at least one buyer can engage in any transactions that change their collective wealth. A coalition that contains player 1 is worth the maximum that the object in question is worth to any member of the coalition.

The tools of cooperative game theory applied to this model reflect some of the important features of such an interaction. For example, the core of the game [which for TU games is equal to the set of payoff distributions with the property that the sum of the payoffs to the members of each coalition S is at least $v(S)$] corresponds to the set of outcomes at which the seller sells to the buyer with the higher reservation price, at some price between twenty and thirty dollars, and no other transfers are made. This corresponds to what we would expect if the buyers compete with each other in an auction, for example. Von Neumann and Morgenstern proposed a more comprehensive kind of “solution” for such a game, which today is called a *stable set* or a *von Neumann–Morgenstern solution*. There are infinitely many von Neumann–Morgenstern solutions to this game, each of which consists of the core plus a continuous curve corresponding to a rule for sharing between the two buyers the wealth at each price less than twenty dollars (should they be able to agree to avoid bidding against one another, for example).

Von Neumann and Morgenstern’s interpretation of this multiplicity of solutions was that each represented a particular “standard of behavior” that might be exhibited by rational players of the game. Which standard of behavior we might expect to observe in a particular game would generally

Cambridge University Press

0521021332 - The Shapley Value: Essays in Honor of Lloyd S. Shapley

Edited by Alvin E. Roth

Excerpt

[More information](#)**4 Alvin E. Roth**

depend on features of the environment – for example, institutional, social, or historical features – not modeled by the characteristic function. Thus their view was that much of the complexity of strategic interactions that was omitted from the characteristic function model reemerged through the complexity of the set of solutions. This very complexity nevertheless made it difficult to make a simple evaluation of a game in terms of its von Neumann–Morgenstern solutions. Partly for this reason, much of the subsequent analysis of such games has focused instead on the core. Although the core is much simpler than the von Neumann–Morgenstern solutions, it may be empty in some games and a large set of outcomes in others. And various “noncompetitive” modes of behavior (such as the formation of a bidders’ cartel in our earlier example) might lead to outcomes outside the core, so a great deal of complexity remains.

This complexity is to a large extent a reflection of the underlying complexity of strategic interaction. Indeed, much current work in game theory is in the direction of putting more institutional and other detail into game-theoretic models in order to be able to more fully describe and better understand these complexities. (To a certain extent the same can be said of individual choice theory, in which there has been in recent years some exploration of more complex models than utility maximization.) However the underlying complexity of the phenomena only increases the need for a simple way to make a preliminary evaluation of games.

1 The Shapley value

Shapley’s 1953 paper (reprinted as Chapter 2 of this volume) proposed to fill this need, essentially by carrying the reductionist program of von Neumann and Morgenstern a step further. Because it had proved so useful to represent each alternative facing a player by a single number expressing its expected utility, and to summarize the opportunities facing a coalition in a game by a single number expressing its worth in units of transferable utility, Shapley proposed to summarize the complex possibilities facing each player in a game in characteristic function form by a single number representing the “value” of playing the game. Thus the value of a game with a set $N = \{1, \dots, n\}$ of players would be a vector of n numbers representing the value of playing the game in each of its n positions. The connection to what I have called the reductionist program of von Neumann and Morgenstern is made clearly in the first paragraph of Shapley’s paper, which begins “At the foundation of the theory of

Cambridge University Press

0521021332 - The Shapley Value: Essays in Honor of Lloyd S. Shapley

Edited by Alvin E. Roth

Excerpt

[More information](#)

Introduction to the Shapley value

5

games is the assumption that the players of a game can evaluate, in their utility scales, every 'prospect' that might arise as a result of a play. . . . [O]ne would normally expect to be permitted to include, in the class of 'prospects,' the prospect of having to play a game."

Shapley's approach was to consider the space of all games that might be played by some potentially very large set of players (denoted by the letter U , to signify the universe of all possible players). In a particular game v , the players actually involved are contained in any *carrier*, which is a subset N of U such that $v(S) = v(S \cap N)$ for any subset of players $S \subset U$. If a carrier N for a game v does not contain some player i , then i is a *null* player, because i does not influence the worth $v(S)$ of any coalition S . So any set containing a carrier is itself a carrier of a game, and any player not contained in every carrier is a null player.

Shapley defined a *value* for games to be a function that assigns to each game v a number $\phi_i(v)$ for each i in U . He proposed that such a function obey three axioms. The *symmetry* axiom requires that the names of the players play no role in determining the value, which should be sensitive only to how the characteristic function responds to the presence of a player in a coalition. In particular, the symmetry axiom requires that players who are treated identically by the characteristic function be treated identically by the value.

The second axiom, usually called the *carrier* axiom, requires that the sum of $\phi_i(v)$ over all players i in any carrier N equal $v(N)$. Because this must hold for *any* carrier, it implies that $\phi_i(v) = 0$ if i is a null player in v . Sometimes this axiom is thought of as consisting of two parts: the *efficiency* axiom ($\sum_{i \in N} \phi_i(v) = v(N)$ for some carrier N), and the *null player* (or sometimes "dummy player"¹) axiom.

The third axiom, now called the *additivity* axiom, requires that, for any games v and w , $\phi(v) + \phi(w) = \phi(v + w)$ (i.e., $\phi_i(v) + \phi_i(w) = \phi_i(v + w)$ for all i in U , where the game $[v + w]$ is defined by $[v + w](S) = v(S) + w(S)$ for any coalition S). This axiom, which specifies how the values of different games must be related to one another, is the driving force behind Shapley's demonstration that there is a unique function ϕ defined on the space of all games that satisfies these three axioms.

The easiest way to understand why this function exists and is unique is to think of a characteristic function v as a vector with $2^U - 1$ components, one for each nonempty subset of U . (For simplicity, take the universe U of players to be finite.) Then the set G of all (not necessarily superadditive) characteristic function games coincides with euclidean space of dimen-

6 Alvin E. Roth

sion $2^U - 1$. The additivity axiom says that if we know a value function on some set of games that constitute an additive basis for G , then we can determine the value for any game.

A set of games that will permit us to accomplish this is the set consisting of the games v_R , defined for each subset R of U by

$$v_R(S) = 1 \quad \text{if } R \subset S, \\ = 0 \quad \text{otherwise.}$$

Any player not in R is a null player in this game, which is sometimes called the *pure bargaining* or *unanimity* game among the players in R , because they must all agree among themselves how to split the available wealth. Because the players in R are all symmetric, the symmetry axiom requires that $\phi_i(v_R) = \phi_j(v_R)$ for all i and j in R . Because the null player axiom requires that $\phi_k(v_R) = 0$ for all k not in R , the efficiency axiom allows us to conclude that $\phi_i(v_R) = 1/r$ for all i in R , where r is the number of players in R . (For any finite coalition S , we will denote by s the number of players in S .) Thus the value is uniquely defined on all games of the form v_R or, for that matter, on games of the form cv_R for any number c (where $cv_R(S) = c$ if $R \subset S$ and 0 otherwise). (Note that cv_R is superadditive when c is non-negative.)

But the games v_R form a basis for the set of all games, because there are $2^U - 1$ of them, one for each nonempty subset R of U , and because they are linearly independent. Therefore any game v can be written as the sum of games of the form cv_R . (For example, the game Γ_1 discussed earlier with one seller and two buyers is given by $\Gamma_1 = 10v_{(1)} + 10v_{(12)} + 20v_{(13)} - 10v_{(123)}$.) And so the additivity axiom implies that there is a unique value obeying Shapley's axioms defined on the space of all games.

Shapley showed that this unique value ϕ is

$$\phi_i(v) = \sum_{S \subset N} \frac{(s-1)!(n-s)!}{n!} [v(S) - v(S-i)],$$

where N is any finite carrier of v , with $|N| = n$. This formula expresses the Shapley value for player i in game v as a weighted sum of terms of the form $[v(S) - v(S-i)]$, which are player i 's marginal contribution to coalitions S . (In Chapter 17 Peyton Young shows how the Shapley value may be axiomatized in terms of the marginal contributions.) In fact, $\phi_i(v)$ can be interpreted as the *expected* marginal contribution of player i , where the distribution of coalitions arises in a particular way.

Specifically, suppose the players enter a room in some order and that

Cambridge University Press

0521021332 - The Shapley Value: Essays in Honor of Lloyd S. Shapley

Edited by Alvin E. Roth

Excerpt

[More information](#)

Introduction to the Shapley value

7

all $n!$ orderings of the players in N are equally likely. Then $\phi_i(v)$ is the expected marginal contribution made by player i as she enters the room. To see this, consider any coalition S containing i and observe that the probability that player i enters the room to find precisely the players in $S - i$ already there is $(s - 1)!(n - s)!/n!$. (Out of $n!$ permutations of N there are $(s - 1)!$ different orders in which the first $s - 1$ players can precede i , and $(n - s)!$ different orders in which the remaining $n - s$ players can follow, for a total of $(s - 1)!(n - s)!$ permutations in which precisely the players $S - i$ precede i .)

Although this is not meant to be a literal model of coalition formation (a topic that will be addressed by two of the papers in this volume), thinking of the value in this way is often a useful computational device. In our example of one seller and two buyers, the three players can enter in six possible orders. If they enter in order 1,2,3, their marginal contributions are (0,20,10), and their Shapley values are the average of these marginal contributions over all six orders: $\phi(v) = (18.33, 3.33, 8.33)$.

For a more challenging example, consider a game loosely modeled on the United Nations Security Council, which consists of fifteen members. Five of these are permanent members and have a veto, and ten are rotating members. The voting rule is that a motion is passed if it receives nine votes and no vetoes. We model this here by taking $v(S) = 1$ if S contains all five permanent members and four or more other members, and $v(S) = 0$ otherwise.

Because $15!$ is a number on the order of 10^{12} , we obviously cannot proceed to calculate the Shapley value by enumerating all possible orderings of the players. But we can use the random-order property, together with symmetry and efficiency, to calculate the value. To do this, note that by symmetry all rotating members have the same value $\phi_r(v)$, all permanent members have the same value $\phi_p(v)$, and efficiency requires that $10\phi_r(v) + 5\phi_p(v) = 1$. In order for a rotating member to make a positive marginal contribution in a random order, all five permanent members and exactly three of the other nine rotating members must precede him or her. There are $9!/3!6!$ such coalitions, corresponding to the different ways to choose three out of the nine other rotating members. As we said, each such coalition S (of size $s = 9$) occurs with probability $(s - 1)!(n - s)!/n!$, and the marginal contribution of the last rotating member is $[v(S) - v(S - i)] = 1$. So the Shapley value of a rotating member is $\phi_r(v) = (9!/3!6!)(8!6!/15!) = .00186$, and the Shapley value of a permanent member is $\phi_p(v) = (1 - 10\phi_r(v))/5 = .196$, which is over 100 times greater.

Cambridge University Press

0521021332 - The Shapley Value: Essays in Honor of Lloyd S. Shapley

Edited by Alvin E. Roth

Excerpt

[More information](#)

8 Alvin E. Roth

1.1 *The Shapley–Shubik Index*

The results of a similar calculation, using the then existing rules of the Security Council, are described in the 1954 paper of Shapley and Martin Shubik, which is reprinted as Chapter 3. That paper was the first to propose applying the Shapley value to the class of simple games, which are natural models of voting rules.

A *simple game* is a game represented by a characteristic function v that takes on only the values 0 and 1. A coalition S is called *winning* if $v(S) = 1$, and *losing* if $v(S) = 0$, and the usual assumption is that every coalition containing a winning coalition is itself winning or, equivalently, that every subset of a losing coalition is itself a losing coalition. (A simple game is called *proper* if the complement of a winning coalition is always losing.) If v is a simple game among some set N of players, then an equivalent representation is simply the list of winning coalitions of N , or even the list of minimal winning coalitions (i.e., winning coalitions none of whose subsets are winning). For some classes of games, even terser representations may be natural: For example, a “weighted majority game” with n voters, such as might arise among the stockholders of a corporation, can be represented by the vector $[q; w_1, \dots, w_n]$, where w_i denotes the number of votes cast by player i , and q denotes the number of votes needed by a winning coalition. The winning coalitions are then precisely those coalitions S with enough votes; that is, S is winning if and only if $\sum_{i \in S} w_i \geq q$.

Because simple games are essentially no more than lists of winning coalitions, they are often natural models of situations in which the full weight of the usual assumptions about characteristic function games may not be justified. Thus, for example, we may want to model a bicameral legislature by noting that the winning coalitions are those containing a majority of members in each house, and without assuming that the log-rolling opportunities available to members are sufficient so that a winning coalition can divide up the spoils in a transferable utility way. When we are interpreting a simple game as something other than a transferable utility characteristic function game, we may want to interpret the Shapley value of each player differently than we otherwise would. In recognition of this, the Shapley value applied to simple games is often called the *Shapley–Shubik index*. The marginal values $[v(S) - v(S - i)]$ in a simple game are always equal to 0 or 1, so a player’s Shapley–Shubik index equals the proportion of random orders in which he or she is a “pivotal” player, the proportion of orders in which the set of players $S - i$ who

Introduction to the Shapley value

9

precede him or her from a losing coalition that is transformed into a winning coalition S by the arrival of player i . (In each ordering of the players, only one player is pivotal.)

Analyzing voting rules that are modeled as simple games abstracts away from the particular personalities and political interests present in particular voting environments, but this abstraction is what makes the analysis focus on the rules themselves rather than on other aspects of the political environment. This kind of analysis seems to be just what is needed to analyze the voting rules in a new constitution, for example, long before the specific issues to be voted on arise or the specific factions and personalities that will be involved can be identified.

The task of assessing how much influence a voting system gives to each voter has assumed legal importance in evaluating legislative reapportionment schemes, following court rulings that valid schemes must give voters equal representation (i.e., must be systems that give “one man, one vote”). This has proved a difficult concept to define when voters are represented by legislators elected by district, particularly in systems in which districts of different sizes may be represented by different numbers of legislators or by legislators with different numbers of votes. A measure of voter influence related to the Shapley–Shubik index, called the Banzhaf index after the lawyer who formulated it in this context (Banzhaf 1965, 1968; Coleman 1971; Shapley 1977), has gained a measure of legal authority, particularly in New York State, in court decisions concerning these issues (Lucas 1983). Instead of looking at random orders of players, the Banzhaf index simply counts the number of coalitions in which a player is a “swing” voter. That is, the Banzhaf index of a voter i is proportional to the number of coalitions S such that S is winning but $S - i$ is losing. (A comprehensive treatment of the mathematical properties of the Banzhaf index is given by Dubey and Shapley 1979). Although the Banzhaf and Shapley–Shubik indices have certain obvious similarities, in any particular game they may not only give different numerical evaluations of a player’s position but they may rank players differently, so the voter with more influence according to the Shapley–Shubik index may have less influence according to the Banzhaf index.

2 The other papers in this volume

Chapters 2 and 3, by Shapley and by Shapley and Shubik, are the “ancestral” papers from which the rest of the papers in this volume follow. Chapters 4 through 10 are concerned with reformulating these ideas in

Cambridge University Press

0521021332 - The Shapley Value: Essays in Honor of Lloyd S. Shapley

Edited by Alvin E. Roth

Excerpt

[More information](#)

10 Alvin E. Roth

order to better understand them. As often as not, these reformulations also lead to generalizations, so by coming to understand the Shapley value or Shapley–Shubik index in new ways, we are also led to different ways to assess the value of playing a game or of measuring the influence of a voter.

2.1 *Reformulations and generalizations*

Chapter 4, “The expected utility of playing a game,” investigates the implications of taking seriously the idea that the Shapley value can be interpreted as a utility function. It turns out that there is a strong and precise analogy between the Shapley value as a utility for positions in games and the expected value as a utility for monetary gambles, because both are risk-neutral utility functions. However, two kinds of risk neutrality are involved in interpreting the Shapley value as a utility: The first involves gambles (“ordinary risk”) among games, and the second involves games that need not involve any probabilistic uncertainty but only the *strategic* risk associated with the unknown outcome of the interactions among the players. Neutrality to ordinary risk turns out to be equivalent to additivity of the utility function, and neutrality to strategic risk turns out to be closely associated with the efficiency axiom. The class of utility functions that represent preferences that are not neutral to strategic risk (and that are therefore “inefficient” value functions) is also characterized, provided that the preferences remain neutral to ordinary risk over games. (The characterization of utilities for preferences that are not neutral to ordinary risk remains an open problem.) The chapter concludes by considering the implications of this for understanding the comparisons among positions in games that are implicit in the Shapley value.

Philip Straffin’s chapter, “The Shapley–Shubik and Banzhaf power indices as probabilities,” is concerned with simple games, and shows that both the Shapley–Shubik and Banzhaf indices can be interpreted as the answer to the question: “What is the probability that a given voter’s vote will affect the outcome of the vote on a bill?” To pose this question, one needs to specify a model of voter probabilities. Straffin observes that the Shapley–Shubik index answers this question if we assume voters’ opinions are homogeneous in a certain sense, and the Banzhaf index gives the answer if we assume voters’ opinions are independent in a particular way. His analysis not only casts new light on the similarities and differences between these two indices, but also suggests how this method of modeling voters might be adapted to particular situations to create new indices when other assumptions about voters are appropriate.