

Contest Analysis Applied to Game Theory

This article shows how a new technique, based upon the *contest analysis algorithm* can be used to solve game-theory problems. By means of this special algorithm, the method solves every problem--even removing the dilemma from the famous Prisoner's Dilemma. It uses the algorithm to determine *forces* among game participants, based upon payoffs. The forces can then be used to compute numbers which show the quality for each single-play outcome in the game. The whole-game mix of these outcomes is determined by the probabilities used for the players' strategies, and the quality of the whole game can be maximized with proper choice of these probabilities. This maximizing set of probabilities is the solution of the game. The game solution makes the game as fair as the potential payoffs (and the forces they generate) allow it to be.

The algorithm has a long history of use in simulations of military combat, but it is new to the field of social choice and game theory. Until now, no general scheme for solving all game theory problems has been found, and this research corrects that deficiency. (Contest analysis makes up the entire content of this Web site. A thorough discussion of its other aspects can be found by following the links at the end of this report.)

What is game theory?

What I know about game theory was gleaned from an excellent book by Philip D. Straffin called *Game Theory and Strategy*. I shall refer frequently to that book, using 'Straffin' to designate the reference. Straffin's description of what is meant by game theory is quoted here:

"Game theory is the logical analysis of situations of conflict and cooperation. More specifically, a *game* is defined to be any situation in which

- i) There are at least two *players*. A player may be an individual, but it may also be a more general entity like a company, a nation, or even a biological species.
- ii) Each player has a number of possible *strategies*, courses of action which he or she may choose to follow.
- iii) The strategies chosen by each player determine the *outcome* of the game.
- iv) Associated to each possible outcome of the game is a collection of numerical *payoffs*, one to each player. These payoffs represent the value of the outcome to the different players....

"...*Game theory* is the study of how players should rationally play games. Each player would like the game to end in an outcome which gives him as large a payoff as possible."

The payoffs are often called *utilities*, especially when the numbers are intangible or subjective (i.e., not measurable quantities like money, time, or goods), and they can be positive or negative. A game is called *zero sum* if the payoffs in each outcome total zero (as if the players make payments to each other after each play of the game).

If each player chooses from among his strategies in a probabilistic way (according to a set of probabilities) while repeatedly playing the game, there will be a long-run average outcome for each

player. A solution to the game is the determination of a set of strategy probabilities for each player which is optimal--giving the best long-run average outcome for each player which is consistent with fairness in the game. The sticking point here is in deciding what is fair. A so-called equilibrium solution (when it exists and is unique) is often taken as the 'true' solution because it has a certain stability (unless good reasons exist for choosing another--such as its being better for all players). When there is no unique equilibrium solution, other methods are used, generally disagreeing with each other and leading to debate. The search has gone on for over 50 years and no all-purpose method has been found until now. (I have the impertinence to claim in this report that a certain implementation of the contest analysis algorithm is an all-purpose method.)

The contest analysis approach to game theory

The zero-sum game always has a numerical solution, but the solution is often unsatisfactory, requiring subsequent rationalization. The situation in non-zero-sum games is far worse, and I quote Straffin again to bring this out: "... the solution theory for zero-sum games does not carry over to non-zero-sum games, and in fact ... there is no cogent general solution concept for non-zero-sum games. We simply cannot give a general prescription for how to play all such games"

It is my purpose in this article to show that the contest analysis algorithm does give a "general prescription" for all games, since it works every time. The problem is set up as a two-sided contest--two teams competing--with all of the *strategies* making up one team and all of the *outcomes* making up the other team. It is a contest of strategies versus outcomes. A vector space of many dimensions is used; every strategy and every outcome has its own direction in this space, and interactions between strategies and outcomes are shown as vectors. (I am using 'outcome' here to mean the outcome of a single play of the game, in which the players each use just one strategy. Sometimes I will use the word to mean the average outcome of an entire game, where the players are considered to play an infinite number of times using a mix of their strategies. The distinction between these two meanings should be clear from the context.)

The data for contest analysis are always in the form of a *scoreboard*. In game theory, this scoreboard is a data matrix with one section having its rows representing strategies and columns representing outcomes and another section with row and column representations reversed. The input data elements are the expected payoffs; converting them to preference distribution functions is the first step of the analysis. The whole sequence of steps will be explained later by means of examples.

The goal for each player in a game is to seek his greatest long-term payoff in repeated plays of the game, with his strategy probabilities somehow in balance with those of the other players. When the balance of strategies feels right, we declare it a fair game and say we have a solution. The vector approach of contest analysis allows us to compute a numerical quantity called *quality*--for each single-play outcome and for the game as a whole--and this number includes both the sizes and distributions of players' gains. For the game as a whole, the quality depends upon the strategy probabilities and a particular set of these probabilities will produce a maximum in the quality function. That set of probabilities becomes the *maximum-quality solution* of the game (or simply the *solution*, replacing current methods).

This new method uses a nonlinear algorithm, operating on the scoreboard, to compute vectors which show the effect of contest participants upon each other. In the case of the *game*, the participants are not the players, but the players' strategies and the outcomes produced by them. The algorithm determines the forces exerted by the strategies upon the outcomes and the reverse. A strong force will be exerted by a strategy in the direction of an outcome if the outcome's payoff to the player is great and if the strategy is

itself strong. The outcomes exert forces upon the strategies, too, with the force depending upon payoff and strength of the outcome. Each participant (strategy or outcome) has its own direction in a vector space of n dimensions, where n is the number of participants. The vector sum of all forces exerted by a participant is his *strength* vector, and the vector sum of all strength vectors is the *contest* vector. (The magnitude of the strength vector is simply the *strength* of the participant--a scalar. The projection of a strength vector upon the contest vector is also a scalar--called the *applied strength* of a participant.)

In combat analysis, where this method originated, the participants are typically combat vehicles of various types attacking and defending in an engagement and sustaining losses. The data are the fractional losses of vehicles, displayed according to which opposing vehicle type caused the loss (a loss-by-cause scoreboard). A vehicle is simultaneously in the role of attacker and defender--shooter and target. From the array of vehicle-versus-vehicle fractional losses, the contest algorithm finds the strengths of all of the vehicles. Attack and defense imply designation, which means pointing at something or being pointed at by something. In combat, it is easy to visualize pointing--by an attacker at a target and by a target at its attacker. Pointing brings in vectors; more-vigorous pointing goes with stronger vectors. But game theory has its own system of pointing, having to do with choosing or being chosen; it just happens that the opposing participants in game theory--strategies and outcomes--are not trying to destroy each other. When a strategy is chosen, vectors automatically point to outcomes that contain that strategy and from those outcomes back to the original strategy. Since pointing at an outcome or at a strategy can be thought of as pointing in a direction parallel to the axis 'owned' by that outcome or strategy, each single participant will produce two sets of orthogonal sub-vectors: one set for *choosing* and the other set for *chosen by*. These two sets are analogous to attack and defense in combat, so it is convenient in discussing game theory to use the terms attack and defense when referring to choosing or being chosen.

The participants that are *outcomes* make up one side of this contest, as mentioned earlier, and their attack vectors can be added (using vector addition) to produce an *outcome vector* for the contest. This outcome vector will have an orthogonal component in the direction of each strategy; these components, together with the probabilities of play of strategies, will be all that is needed to compute the *quality* of the game. The components of the outcome vector show the rewards given by the game to the strategies--and hence to the players. Having these rewards be both large and fairly distributed to the players is the situation that the game solution seeks--the maximum-quality solution--and there is always a set of strategy probabilities which maximizes the game's quality.

While it is clear enough that the attacker can be pushing toward the defender, the reader may wonder at this point how a participant on the defending side can be exerting a force in the direction of the attacker. For this, I must refer to the mechanical analog used elsewhere on this site in deriving the equations for the algorithm. Briefly, the defender operates a brake which keeps the force of the attacker from causing a body to move. Every defender has a separate brake for each of the attack bodies he must keep stationary; the movable body and its associated brake is the device which connects attacker and defender.

Participant A, on attack, uses an adjustable lever to push a body in the direction of participant B, on defense, who presses a brake pad against the body to keep it stationary. The direction of the brake-pad force must be at right angles to the allowed direction of motion of the body, and we define that brake-pad direction as toward participant A. The same thing happens when attacker/defender roles are reversed. Thus an attacker has a vector made up of components in the directions of those he is attacking, and a defender has a vector made up of components in the directions of his attackers. A participant's *strength* is the *magnitude* of his attack and defense vectors (the same magnitude for attack and defense), and the *directions* of his attack and defense vectors will generally be different. The analogy is a purely

mathematical construct, since it extends the normal rules of mechanics to an imaginary space of any number of dimensions. The objective of the computation in the algorithm is to determine the magnitudes and directions of all of the vectors.

A numerical example to show how the contest-analysis procedure works in game theory

The first example I call the skewed roulette problem: a roulette wheel with only red and black to be played (no green spaces) and with more red spaces than black. Only one player has a choice (red or black for a bet) and the other player is the operator of the wheel, governed by what the wheel produces. So the game is between a player and the wheel. (A one-sided game like this is referred to as a game against nature.) The player bets on either red or black and the wheel must pay the amount of the bet to the player if the wheel's color comes up to match what the player chose. If the wheel's color does not match, the player loses his bet. All bets are of one unit (a chip).

There are four strategies in the problem: red or black for the player and red or black for the wheel. The wheel's strategies have fixed probabilities, but the player can use any probability of play that he wishes. In this example, we will give the wheel 60% red spaces. The player will do best for himself by placing every bet on red--taking full advantage of the wheel's skewness. Maximum fairness in the game (zero expected gain for both sides) will be achieved if he splits his betting 50-50 between red and black. (Only one side needs to have 50-50 splitting and, in normal roulette, the wheel takes care of that part, except for the green spaces.) So how can we specify a probability for betting on red which will give maximum quality to the game?

The strategies will be labeled 1, 2, 3, and 4 with 1 and 2 for the player and the others for the wheel. The payoff matrix looks like this, where the first entry is payoff to the player and the second is payoff to the wheel:

	3	4
1	(1, -1)	(-1, 1)
2	(-1, 1)	(1, -1)

The four matrix locations represent the outcomes, so the payoff matrix can be rewritten in strategy-vs.-outcome form as follows (where the columns now represent the four outcomes):

1:	1	-1	--	--
2:	--	--	-1	1
3:	-1	--	1	--
4:	--	1	--	-1

The payoffs cannot be negative so they are adjusted by adding 1 to all of them and then normalizing to unity. The result is an array of *preference distribution functions* for player and wheel, shown next.

1:	.5	0	--	--
2:	--	--	0	.5
3:	0	--	.5	--
4:	--	.5	--	0

The strategy probabilities are set aside, for the moment, while we form a scoreboard for the contest analysis--strategies vs. outcomes--with scores being preferences. We already have the preferences of strategies for outcomes and we need to complete the rest of the scoreboard by entering the preferences

of outcomes for strategies. The scoreboard is shown below, with the lower-left section (outcomes as attackers) showing how much the outcomes (rows) are rewarded by the strategies, in terms of being preferred. The preference of a strategy for an outcome is really the preference of the given strategy for the combination of *other* strategies in the outcome. This example, with two players, has only one *other* strategy (no combination).

For the upper-right section (strategies as attackers) we need to think of the preference of an outcome for a strategy as the preference that the combination of *other* strategies has for the given strategy. For the two-player example, that is the same as the preference of the *other* strategy for the outcome. (How combinations are treated--games with three or more players--will be discussed later.) Thus the upper-right section of the scoreboard contains the same numbers as the lower-left, but with rows and columns interchanged and with the resulting column entries also interchanged.

A score in the array is a quantity which enhances the strength of the row's participant as attacker. For participant 1 (the first strategy of the first player), his strength is all due to outcome 6 having some preference for him--but that preference in outcome 6 is really the preference of strategy 4 (second strategy of the second player) for strategy 1. One can go through the scoreboard and see how the numbers of the original preference distribution functions get applied as scores.

	1	2	3	4	5	6	7	8
1:	--	--	--	--	0	.5	--	--
2:	--	--	--	--	--	--	.5	0
3:	--	--	--	--	.5	--	0	--
4:	--	--	--	--	--	0	--	.5
5:	.5	--	0	--	--	--	--	--
6:	0	--	--	.5	--	--	--	--
7:	--	0	.5	--	--	--	--	--
8:	--	.5	--	0	--	--	--	--

The next step is to use the contest algorithm to find the vectors of the game for this scoreboard. The algorithm is derived elsewhere, so I will just write the general equation here: $X_i = [\sum_j A_{ij} \cdot f_{ij} \cdot X_j]^{1/2}$ where X_i and X_j are the strengths of participants i and j (attacker and defender, column and row) and A_{ij} is the scoreboard element i,j . The quantity f_{ij} is the fraction $A_{ij} / \sum_i(A_{ij})$, which is A_{ij} divided by the sum of the column containing A_{ij} . The coupled simultaneous equations of this form make up the algorithm, and a numerical solution for the X values of the set is easy to obtain.

From inspection of the scoreboard, we see that all participants must have the same strength, leading to a quick solution of the above equations for the X values: each one is 0.5. The solution vectors for this particular case are shown in the array below, in which the elements are vector components--the component of the strength vector of the row participant which points in the direction of the column participant. It just happens that the numerical values in the solution array turn out to be identical to the scoreboard values in this simple case.

	1	2	3	4	5	6	7	8
1:	--	--	--	--	0	.5	--	--
2:	--	--	--	--	--	--	.5	0
3:	--	--	--	--	.5	--	0	--
4:	--	--	--	--	--	0	--	.5
5:	.5	--	0	--	--	--	--	--
6:	0	--	--	.5	--	--	--	--
7:	--	0	.5	--	--	--	--	--
8:	--	.5	--	0	--	--	--	--

The magnitude of the contest vector is quickly found by squaring the sum of each column, adding up the squares, and taking the square root of the sum. This gives a contest magnitude of square root of 2. Of more interest in game theory is the magnitude of the outcome vector (lower-left vector sum). That one has magnitude 1, for this example.

The next step of the analysis utilizes only the outcome vector and its array (lower left). This time we make an array of scalar elements from the outcome vector's array of vector components. The elements of this scalar array are the projections of the vector components of the previous matrix upon the outcome vector. This projection can be shown with a bit of algebra to have a very simple form: the scoreboard entry multiplied by the strength (X) of the participant whose preference generated the entry and then divided by the magnitude of the outcome vector. This scalar array is shown next.

	1	2	3	4
5:	.25	--	0	--
6:	0	--	--	.25
7:	--	0	.25	--
8:	--	.25	--	0

The sum of all projections is equal to the outcome-vector magnitude, so dividing each element of the array of projections again by the outcome magnitude will give yet another array, a normalized one (in which the elements all add up to unity). Since the outcome magnitude was shown above to be unity, the last array is already normalized. Elements of this last array are the fractional contributions of the strategies to the strength of a potential game outcome in which the *strategies all have the same frequency of use*. The elements are *weighted preferences of strategies for outcomes* (since preference values were weighted by strategy strength in the computation).

I define the *quality* of a single-play outcome as the product of that outcome's row entries; all of these have zero quality, in this example. Quality is thus the *joint weighted preference* of the strategies in the outcome. For an individual outcome, quality comes from the *potential* of being realized in the game, not from how often it is realized in the play. For this game to have a nonzero realized quality, a mix of strategies must be used in such a way that the player and the wheel operator can accumulate their preference scores from one spin to the next. The solution of the game is the specification of the strategy mix which maximizes the realized quality. The next step, then, is to multiply each row of the last matrix by its probability of realization: the joint probability of play of its strategies. Let p_1 , p_2 , p_3 , and p_4 represent the probabilities of play for strategies 1 to 4 and have $p_3 = 0.6$ and $p_4 = 0.4$ (the probabilities for the wheel to come up with red or black, respectively). Then, with $p_2 = (1 - p_1)$, we are left with finding only p_1 to maximize the game fairness.

Rewriting the last matrix with an additional column for the row multiplier (probability of occurrence of the outcome), we get this:

	<u>1</u>	<u>2</u>	<u>3</u>	<u>4</u>	<u>Row multiplier</u>
5:	.25	--	0	--	(p ₁) (.6)
6:	0	--	--	.25	(p ₁) (.4)
7:	--	0	.25	--	(1-p ₁) (.6)
8:	--	.25	--	0	(1-p ₁) (.4)

Multiplying each row by its multiplier and then adding up each column, we get a sum (Y_i) for each of the four columns. Columns 1 and 2 represent the player's strategies, with columns 3 and 4 for the operator of the wheel. The sum (Y₁ + Y₂) is the *expected value* of the weighted preference obtained by the player when the strategy probabilities are as given by the row multiplier. The sum (Y₃ + Y₄) is the same for the wheel. The quality of the game is the product (Y₁ + Y₂)(Y₃ + Y₄), the expected values for player and wheel multiplied together. A little algebra discloses that the maximum quality occurs when p₁ is 0.5, and that is true for any amount of skewness that the wheel may have. The player must use a 50-50 split if he is going to play according to the game solution of maximum quality. The theory will not let the player take advantage of the skewness of the wheel.

Discussions of game theory often make use of the concept of how a logical player would maximize his gain while limiting that of his opponent. Here, that would be to bet on red every time if the wheel is skewed toward red, but the game would probably end quickly when one side realizes what is happening. The aim of game theory should be to specify fair rules for the game and then act as a referee. In roulette, the referee is telling the player that he can either play any way at all against a fair wheel or play with a 50-50 split against a skewed wheel. If he is unsure of the wheel, he can guarantee fairness by using the 50-50 split in his betting, randomized by flipping a coin for play of red or black. (I am talking about a simplified red-black wheel here, without the extra "house" spaces.)

If the question were to decide how to do single play (no repetition), the player must still use his coin flip to make a bet on red or black, giving him an expected payoff of zero but a realized plus or minus payoff.

Digression on use of strategies' preferences for outcomes

In the last example, I had the preference of an outcome for a strategy be the same as the preference of the other player's strategy for the outcome, because the outcome's preferences are really nothing more than the preferences of the two strategies for each other. Then the question arises about how to handle games with three or more persons, when there are two or more 'other' strategies in each outcome. For this explanation, a three-person game will suffice, since extension to higher orders is straightforward. In a three-person game, each outcome will involve three strategies, one from each player. Each player will have preferences (summing to 1) distributed over the strategies. Let us use the symbol p₁ for the preference of a particular strategy of player 1 for a particular outcome. The other players will have preferences p₂ and p₃ for their preferences for the same outcome. The joint preference for the outcome is the product of the three preference quantities, and that product can be rewritten in the following way:

$$p_1 \cdot p_2 \cdot p_3 = (p_1 \cdot p_2)^{1/2} (p_1 \cdot p_3)^{1/2} (p_2 \cdot p_3)^{1/2}$$

Each of the three square-root factors can then be taken to be the preference of a two-player combination of strategies for the excluded strategy. That is, the first factor is the preference of the 1,2 combination for the strategy of player 3 going into this particular outcome. The contest algorithm needs to have the preferences going both ways--strategy for outcome and outcome for strategy--and this logic shows how we can get those numbers. If the number of players is four, the logic is the same, except that there are three 'other' strategies in each outcome, and the starting point is to use 1/3 rather than 1/2 for the

exponent which splits the preferences.

The scoreboard row elements are the preferences of the column participants for the participant of the row. If the row is that of an outcome, the element will be the preference of the strategy for the outcome. If the row is a strategy, the situation is more complicated in that the element will be the combined preference of the other strategies for the first one, as illustrated above.

Example with two players, two strategies for each, and no strategy fixed

This is the old standby, the Prisoner's Dilemma, using the version shown as game 12.1 in Straffin's book. The raw utility data are shown below, where the first number in each location is the payoff to player A and the second number is the payoff to player B. The rows represent the two strategies of player A and the columns the two strategies of player B. This problem gets its name from the situation it could describe in which two prisoners are offered release if neither confesses to a crime (the 0,0 status-quo outcome) and a mild sentence if both confess (the -1,-1 outcome). If one confesses, he gets a reward while the other gets a severe sentence (the off-diagonal outcomes). The Prisoner's Dilemma is much studied because standard game-theory rules point to what appears to be the wrong outcome: The second strategy for both players is better than the first no matter what the other player does--dominates the first strategy--yet the outcome (-1,-1) which results from both players playing their second strategies is worse than if they had both played their weaker first strategies and obtained 0,0.

		B	
		(0, 0)	(-2, 1)
A			
		(1, -2)	(-1, -1)

For the algorithm, the payoff data need to be arrayed as strategies versus outcomes, thus:

A:	0	-2	-	-
	-	-	1	-1
B:	0	-	-2	-
	-	1	-	-1

Converting payoffs to preference distribution functions leads to the following array:

A:	1/3	0	---	---
	---	---	1/2	1/6
B:	1/3	---	0	---
	---	1/2	---	1/6

The scoreboard for the contest algorithm uses these last data and is shown next. The lower-left part is just the transpose of the previous preference array, and the upper-right part is formed by taking the preference array and interchanging the players' entries.

	1	2	3	4	5	6	7	8
1:	---	---	---	---	1/3	1/2	---	---
2:	---	---	---	---	---	---	0	1/6
3:	---	---	---	---	1/3	---	1/2	---
4:	---	---	---	---	---	0	---	1/6
5:	1/3	---	1/3	---	---	---	---	---
6:	0	---	---	1/2	---	---	---	---
7:	---	1/2	0	---	---	---	---	---
8:	---	1/6	---	1/6	---	---	---	---

Application of the contest algorithm gives strengths of the participants. Only the strengths of the strategies are needed for the rest of the analysis and they are as follows, for X_1 through X_4 : 0.4202, 0.0833, 0.4202, and 0.0833. The strength of X_i will multiply the all column- i elements of the previous matrix and the result for the lower-left section is shown below. (The elements do not sum to unity because I did not normalize this time.) The column called 'row multiplier' contains the product of the two strategy probabilities involved in each outcome, which product is the probability for occurrence of the outcome in the row.

	<u>1</u>	<u>2</u>	<u>3</u>	<u>4</u>	<u>Row multiplier</u>
5:	.140	---	.140	---	$(p_1)(p_3)$
6:	0	---	---	.042	$(p_1)(1-p_3)$
7:	---	.042	0	---	$(1-p_1)(p_3)$
8:	---	.014	---	.014	$(1-p_1)(1-p_3)$

It is quite clear from inspection that outcome 5 (both components large) is likely to give the maximum-quality solution, and variation of the strategy probabilities confirms this. First the p_1 and p_3 values are picked and the row multiplier is used to multiply the matrix elements of the row. Next, the modified columns are summed and the sum of the first two columns is multiplied by the sum of the second two columns. (We are obtaining the product of the *expected weighted preferences* of the two players for outcomes.) This product is the game quality, and the probabilities which maximize the quality specify the game solution--which occurs when p_1 and p_3 are both 1. Both players should use their first strategies exclusively, producing only outcome 5.

The payoff to both players is zero. The solution arrived at by the conventional method is for complete use of outcome 8, giving a worse payoff for both players--and giving rise to many years of discussion about the peculiarity of the Prisoner's Dilemma. When the problem is solved with this vector method, there is nothing at all peculiar about it.

I have not been able to adjust inputs, while keeping it as a Prisoner's Dilemma problem, and obtain any other maximum-quality solution than the pure-strategy one given above.

Some more worked examples

In this section I will present some more examples from Straffin--only the inputs and outputs, with some discussion. The first is a zero-sum game with an equilibrium point, Straffin's example 2.1. This is a two-player game with four strategies per player. Since the input data array is too big to fit here in its template form, I will just show the data in the same matrix form as in Straffin's book (rows for player A and columns for player B). This is a zero-sum game, but I will put both payoffs in each cell, even though one is the negative of the other.

	<u>B</u>			
A	12, -12	-1, 1	1, -1	0, 0
	5, -5	1, -1	7, -7	-20, 20
	3, -3	<u>2, -2</u>	4, -4	3, -3
	-16, 16	0, 0	0, 0	16, -16

Straffin shows that this example has an equilibrium at the 2,-2 outcome (strategies A3 and B2, underlined) and traditional theory says that this equilibrium denotes "rational" play and should be the game solution. The quality-maximization procedure gives a mixed-strategy solution, with $p_2 = 1$ and player B using three strategies: $p_5 = 0.64$, $p_7 = 0.05$, and $p_8 = 0.31$. The average payoff is -2.65 and 2.65, respectively for A and B--nearly the reverse of what the conventional (equilibrium-point) solution gives.

On the above payoff array, the solution is about a 2/3 to 1/3 split between the first and fourth elements of row 2, with a small admixture of the third element and no use of the second element.

For single play, player B would use the solution probabilities (random-number draw) to make his strategy choice and produce a payoff for himself of -5, -7, or 20, with the first one most likely. Since both players might object to this large gamble, they might incorporate another rule ahead of time: If the largest individual-outcome quality in the game is greater than that of any outcome in the mix, the players can choose that largest one as the single-play solution. In this example, outcome 23 has this property so the players might go to outcome 23 (A's fourth strategy and B's third) where the outcome is 0,0. Notice that other outcomes have 0,0 as payoffs, but the entire interplay of preferences in the contest selects just one of these. Again, the equilibrium point has no bearing on the solution.

Next, I consider another game from Straffin (game 2.2), similar to the last one but with no equilibrium point. Instead, it has four saddle points. The payoff matrix, with the four saddle points underlined, is as follows:

		B			
	4, -4	<u>2, -2</u>	5, -5	<u>2, -2</u>	
A	2, -2	1, -1	-1, 1	-20, 20	
	3, -3	<u>2, -2</u>	4, -4	<u>2, -2</u>	
	-16, 16	0, 0	16, -16	1, -1	

Saddle points always have equal payoffs to the players and standard game theory has it that the game solution must be at one of the saddle points, any one of which can be taken as the game solution, with all payoffs the same. Using the contest-analysis approach leads to a pure-strategy solution--player A using only his second strategy and player B using his third, the payoffs being -1,1. That solution happens to lie right in the middle of the four saddle points, but I do not know if that will prove to be a general rule. The maximum-quality solution has the payoffs closer together than in the solution found with standard game theory, and reversed in direction.

The single-play solution has to be the same as the repeated-play solution, for this problem, because we have a pure-strategy solution..

Another example will be Straffin's game 11.2, which he deems intractable. This is a two-person game with two strategies for each player. His payoff data placed into my template array are the following:

A:	2	1	-	-
	-	-	3	0
B:	4	-	1	-
	-	0	-	4

Converting the raw data to preferences gives the next array:

A:	1/3	1/6	-	-
	-	-	1/2	0
B:	4/9	-	1/9	-
	-	0	-	4/9

This new form of game theory gives a pure-strategy solution: both A and B using their first strategies, with respective payoffs of 2 and 4. Straffin gives the traditional solution (Nash equilibrium) as probabilities 3/7, 4/7 for player A and 1/2, 1/2 for player B (payoffs 1.50, 2.29) and discusses the reason for that solution being unsatisfactory. The reason has to do with each player ignoring his own payoffs

and violating the so-called Pareto principle, which says that no solution should be accepted if another one gives a better reward for both players (which the maximum-quality solution does). The Pareto principle is not needed for this new way of solving game-theory problems.

Inspection of the preference array shows the quality of the second and fourth outcome to be zero (columns 2 and 4) and the quality of the third to be weak due to low symmetry, leading one to expect the first outcome to be the solution. Only the run of the contest algorithm shows that this is indeed the case, however, since the strengths of the participants as well as their payoffs are involved in the problem.

Next, I consider a three-person Prisoner's Dilemma problem to show what happens in a heretofore unmanageable situation. This example is game 21.1 in Straffin, and I am going to display his payoff data with my template technique, as follows:

A:	1	0	0	-2	-	-	-	-
	-	-	-	-	3	2	2	-1
B:	1	0	-	-	0	-2	-	-
	-	-	3	2	-	-	2	-1
C:	1	-	0	-	0	-	-2	-
	-	3	-	2	-	2	-	-1

The rows represent the strategies for players A, B, and C--strategies A1, A2, B1, B2, C1, and C2--so the template arrangement gives the strategy and outcome placement of all input data. All players are identical due to symmetry. Converting to preference distributions, we have the next array:

A:	1/7	2/21	2/21	0	-	-	-	-
	-	-	-	-	5/21	4/21	4/21	1/21
B:	1/7	2/21	-	-	2/21	0	-	-
	-	-	5/21	4/21	-	-	4/21	1/21
C:	1/7	-	2/21	-	2/21	-	0	-
	-	5/21	-	4/21	-	4/21	-	1/21

This array can be converted to the array of preferences *by outcomes for* strategies, using the logic discussed above, namely having it be the square root of the joint preference of the 'other' two strategies for the outcome. The result is shown below, and this array becomes the upper-right section of the scoreboard for the contest algorithm. Notice that the product of the entries in each column is the combined preference of the strategies for the outcome and that this product is the same in both of these arrays. There is little preference among the outcomes for the second strategy of the players.

A:	.143	.151	.151	.190	-	-	-	-
	-	-	-	-	.095	0	0	.048
B:	.143	.151	-	-	.151	.190	-	-
	-	-	.095	0	-	-	0	.048
C:	.143	-	.151	-	.151	-	.190	-
	-	.095	-	0	-	0	-	.048

The transpose of the preference-distribution array (the first of the above two arrays) becomes the lower-left section of the scoreboard--the preferences *by strategies for* outcomes. I will show this below, with the fractional entries replaced by three-place decimals. I also indexed the rows and columns to show that participants 1-6 are the strategies and 7-14 are the outcomes. Notice that outcome 7 already looks like a winner, due to symmetry, but maybe others could contribute in a repeated-play game (when preferences in one outcome can combine with those in another).

	1	2	3	4	5	6
7:	.143	-	.143	-	.143	-
8:	.095	-	.095	-	-	.238
9:	.095	-	-	.238	.095	-
10:	0	-	-	.190	-	.190
11:	-	.238	.095	-	.095	-
12:	-	.190	0	-	-	.190
13:	-	.190	-	.190	0	-
14:	-	.048	-	.048	-	.048

Running the contest algorithm on the full 14 x 14 scoreboard gives the strengths of all participants, and multiplying the preferences of the last array by the strengths of the column participants gives the array of quantities to be used for quality maximization, as in the previous examples. When this is done, the game solution does indeed turn out to be with all players using their first strategy (outcome 7), with payoff to each being +1.

So again, as with the two-person Prisoner's Dilemma problem above, the solution is not the one arrived at by conventional theory using dominance and equilibrium. (That one would be for the players to use only their second strategies.) Varying inputs, while keeping within the bounds of a Prisoner's Dilemma game, can change the solution. In particular, with the maximum payoffs of all three players changed from 3 to 20, the solution calls for the three to use their first strategy only 81% of the time and get a payoff of 3.08 for each. For a single-play game, they may not want to gamble with their random selection of strategy and may instead pick the best pure-strategy outcome, which is number 7 (giving a payoff of 1 to each).

Straffin goes on to discuss the failure of a coalition approach and the apparent general hopelessness connected with this problem. He shows that a coalition of two players, in conventional game theory, would leave them worse off than the player outside the coalition. Even the grand coalition of all three players would be worse for a player than playing alone. (Coalitions are formed here by averaging the payoffs. I did not yet investigate unequal division of payoffs among coalition members.) But this hopelessness does not occur with the contest-analysis approach, at least for the present example with the original payoffs. The maximum-quality solution, with a two-player coalition, is for all players to use their first strategy all of the time. The payoffs then are the same as they were without the coalition. For the grand coalition, the solution again calls for all players to use their first strategy, with payoffs again of 1 to all players.

An example of a game against nature

Straffin has an interesting example from real life, which he calls *Jamaican fishing*. The data for the problem became game 4.1 in his book. It concerns catching fish in traps which are placed one day and retrieved the next day. (The bad consequences of fishing with traps was pointed out by Jacques Cousteau: Lost or abandoned traps catch fish forever and deplete the stocks. This observation is beside the point here.) Fishermen have a choice of fishing at two locations (inside banks or outside banks) and they have three strategies: inside banks, outside banks, or a split between the two. Nature gives two conditions which complicate the decision for the fishermen: A current can run at random times, and this current (when it runs) makes the fishing somewhat better at the inside location and much worse at the outside location. The data--relative profits in the six conditions of fishing-- are shown below. The negative profit indicates a condition where the fish catch cannot make up for expenses.

		<u>Current</u>	
		<u>Run</u>	<u>Not run</u>
Fishermen	Inside	17.3	11.5
	Outside	-4.4	20.6
	In-Out	5.2	17.0

Straffin is presenting the results of another worker here, and that worker treated this as a zero-sum game and solved it. However, one of the players is nature, represented by the current running or not running, and it cannot be assumed that the payoffs to nature will be simply the negatives of the payoffs to the fishermen (the condition for a zero-sum game). In particular, nature is certainly not concerned with the operating expenses of the fishermen (which figure in the computation of the payoffs). This has to be a non-zero-sum game, and we need to do some thinking about what the payoffs to nature should be. In my opinion, the payoffs to nature should reflect either indifference or a penalty for loss of fish, the former assuming an infinite supply of fish and the latter assuming a limited supply.

The original worker on this problem solved for the probability of the current run along with the probabilities for the strategies of the fishermen, and it was subsequently pointed out that this was a mistake: Nature is not a reasoning entity, and the *observed* probability of current run must be used in any analysis. This probability is 0.25. An expected-value solution of the problem, using the fixed probability for run of the current, admits only one strategy for the fishermen (the one with the maximum expected value): use of the outside location only with a payoff of 14.35. (This computation is in Straffin.) In actual practice, the fishermen are said to avoid that location. This expected-value solution is not treating the problem as a *game*, since it ignores the payoffs to nature.

When nature is assumed to be indifferent, we give all of nature's payoffs the same positive number so that nature's distribution function over the outcomes is flat. The current probabilities are held fixed at the observed values of 25% run and 75% not run. Then the problem has the following input data:

		<u>Current</u>	
		<u>Run (B1)</u>	<u>Not run (B2)</u>
Fishermen	Inside (A1)	(17.3, 1)	(11.5, 1)
	Outside (A2)	(-4.4, 1)	(20.6, 1)
	In-Out (A3)	(5.2, 1)	(17.0, 1)

The maximum-quality solution is for the fishing to be all at the outside location (A2), with a payoff of 14.35 for the fishermen.

When it is assumed that nature wants to conserve the fish stocks, all of nature's utilities will be negative or zero, corresponding to catches of fish in the six fishing situations (inside banks with running current, inside without current, etc.). The bigger the catch, the worse it is for nature. We are not given any real data here for the fish catches, so I will have to work this problem with synthetic data. The next table makes a non-zero-sum problem out of the former zero-sum one by putting in data for player B (nature). I am just saying that the fish catch goes along somewhat in parallel with the profits, allowing for extra expense in fishing outside and even more expense there when the current is running. In reality, more than just the weight of fish would have to be considered, since different types would be caught at the different locations, and some types would be more valuable than others or more deserving of protection.

		<u>Current</u>	
		<u>Run (B1)</u>	<u>Not run (B2)</u>
Fishermen	Inside (A1)	(17.3, -19.3)	(11.5, -13.5)
	Outside (A2)	(-4.4, -0.6)	(20.6, -24.6)
	In-Out (A3)	(5.2, -8.7)	(17.0, -20.0)

Solving this problem with the quality-maximizing procedure, again holding the current probabilities fixed at 25% and 75%, gives the result that the fishing should be all inside (A1), with average payoff of 12.95 for the fishermen and -14.95 for nature. Using the made-up cost data of this problem, the fish-conserving play results in a penalty for the fishermen and a benefit for nature. There is a complete reversal of fishing tactics from the nature-indifferent case. (The data from the field show that none of the fishermen use strategy A2 and most of them use strategy A1.)

I will not try to compare my solutions with those of other workers, since I was just using the problem as an example of a real-world situation that can be nicely treated with this new version of game theory. The example does show that game theory can be thought of as a generator of a set of rules for fair play, in this case preserving fish stocks.

Some discussion about the Prisoner's Dilemma

The general form for Prisoner's Dilemma can be shown with the array below, where letters are used to represent payoffs of different sizes. Player A has payoffs L, ML, MH, and H--standing for low, medium-low, medium-high, and high payoffs. Player B has lower-case versions of the same letters for his payoffs.

		<u>Player B</u>	
<u>Player A</u>		MH, mh	L, h
		H, l	ML, ml

The second strategy of player A is dominant for him, since his second row shows a bigger payoff than the first row. Similarly, the second strategy of player B is dominant for him, since B's second column always gives a bigger payoff than his first column. "Rational" play would be for each to use the dominant strategy all the time, with the outcome at the lower right. But if both play their *other* strategy, they come out at the upper left, with bigger payoffs than at the lower right. (The lower-right outcome is what is known as a Nash equilibrium--the only one in the problem. More will be said about equilibria later; here I will merely state that game theory traditionally looks for solutions from among the equilibria.)

I gather that there has been much thinking and debate in the community of game-theory specialists over the question of which outcome--upper-left or lower-right--should be considered the proper one, as if it has to be one or the other. Straffin thinks that it should be the upper-left one, and he invokes an *ad hoc* rule (after Pareto) which says that the equilibrium outcome should be trumped by another outcome if that other outcome has a better payoff for both players. To the contrary, others give the only choice for the solution as the "rational" equilibrium outcome at the lower right of the array.

With quality maximization, the Prisoner's Dilemma problem has nothing special to distinguish it from any other game-theory problem. All of the fretting about it over the years has come about because everybody wanted to preserve the concepts of dominance and equilibrium for solving problems. My position is that both of those concepts are useless.

The logic behind the whole process

The traditional method of solving problems in game theory has the goal of maximizing payoffs for players, while maintaining fairness--meaning not letting one player gain too much over another. Since

fairness is in the eye of the beholder, the natural recourse is to look for some sort of mathematical optimization--let the logic of mathematics decide the fairness question. For zero-sum, two-player games, a particularly easy form of optimization is available in the *equilibrium* solution. This is a situation in which each player can limit the winnings of his opponent and it is always possible to find an equilibrium solution with either pure strategies (the same strategy played over and over) or mixed strategies (a probabilistic selection of strategy in each repetition of a game). When the two-player game is not zero-sum, the equilibrium has a complicated structure and is given the special name *Nash equilibrium*. What complicates it is that it makes use of extra data that are not in the problem's data set to break the problem into two zero-sum problems: Player A assumes that player B is assuming that player A is opposing him (player B) in a zero-sum game (with negatives of player B's payoffs assumed for player A). This is a lot of assumptions, but the solution of this assumed zero-sum game can always be obtained (with either a pure or mixed strategy). Player A will then take the strategy just found (in what can be called player B's zero-sum game) as his side of the Nash equilibrium. When player B does the same thing with the other zero-sum game (player A's game) he will have found his side of the Nash equilibrium. The short-form description for all this is that each player is responding rationally to what he assumes is rational play on the part of his opponent. It seems to me that these assumptions of zero-sum games do not necessarily constitute rationality, and the introduction of extra data (the missing side of each zero-sum game) contaminates the problem.

The situation for the non-zero-sum problem is that the Nash equilibrium may not offer a solution at all, and for these reasons: It may not be unique, or it may be a solution which is *non-Pareto-optimal* (another solution giving a better payoff for both players), or it may simply not feel right. Or (and this is a big one) the problem may have more than two players, so there is no Nash equilibrium. Other ways of attacking game-theory problems are available, but they all have drawbacks, and their very number is a generator of multiple candidates for solutions.

Even the zero-sum game, with its guaranteed equilibrium solution, has a certain philosophical question about aggressive versus cautious play. The equilibrium is the point at which every player is being the most cautious (or defensive) in his play. For a pure-strategy equilibrium, he is taking the strategy whose minimum payoff (to him) is the greatest. (A mixed-strategy equilibrium solution has something similar in its reasoning.) That may seem to be a fair solution, but maybe it is tilted too far in the direction of defense. Each player is purely defensive when he acts only to limit the payoffs obtainable by his opponents (which is what he is doing when he picks a strategy with the greatest minimum payoff for himself--or lowest maximum for his opponent). One gets the feeling that equilibrium solutions are used because they offer a first crack at a game problem and, if they fail, something else can be tried.

Standard practice in game theory is to look for and eliminate dominated strategies before doing anything else. For certain data sets, strategies can be eliminated sequentially until only one outcome is left, that then being the solution to the problem (in the traditional view). But eliminating strategies also eliminates data and it is generally considered bad form to discard data and then claim to have properly solved a problem. In the Prisoner's Dilemma and other problems, it is amply demonstrated that this practice can lead to a bad solution. With this new method, no strategies are eliminated. In fact, a dominated strategy has a role to play. If such a basic concept as dominance cannot be relied upon, what else in traditional game theory can be trusted? I would say that the game-theory discipline is in trouble if a single method cannot be used for all problems.

The first step in developing an all-purpose method is to use preference numbers to judge the level of agreement among individuals about how much they want something. This points to the need for normalized preference distribution functions on the outcomes. How else could one judge whether the

desires for outcomes are close together? With the contest-analysis approach, we address this point at the very beginning by converting the given utilities of a problem to preference distribution functions. Other applications of contest analysis (pair selection and voting, for example) also use normalized preference distribution functions. Converting a game's utilities to preferences need not be done in the linear way shown in the examples of this report, and converting may require an algorithm of its own. (The utilities may be multidimensional--maybe expressed in terms of both time and money--or the players may feel that a nonlinear conversion better suits their preferences.) In any case, the preference distribution function must be the foundation for game theory. I surely am not the first to recognize the need for preference distribution functions (cardinal preferences), and I take no credit for this.

The second step is to notice that the preference numbers can serve very well as scores in a contest between strategies and outcomes. The contest algorithm was designed to work in situations like this. The probabilities of play do not come into the analysis until later, after the basic contest has been solved. The basic contest treats all outcomes as potential ones, and they all are used in the strength-determining phase of the problem, even if they do not enter into the final solution.

To get a feeling for the contest analysis algorithm itself, and how it is used in a wide variety of applications, the reader is invited to examine the other reports on this site, beginning with the links shown at the end of this report.

The one-shot game

In many applications, there is no question of repeated play of a game and therefore no chance for a mix of strategies to be employed. But single play is still a game and the element of chance is a feature of any game. If the game were to have repeated play and if it has a mixed-strategy solution, random selection of strategies would be expected, beginning with the first play. A single-play game can then be thought of as an infinite-play game which has been stopped after the first round. The players should use their random-number draws to select their strategies even in a one-shot game. But it can happen that the chance aspect can introduce more risk than the players can tolerate in a one-shot game--no later play to recover from a big loss. In that case, the players can agree to take the individual outcome with the highest quality and eliminate the element of chance. This outcome may not be in the original solution mix.

The players could be invited to negotiate among themselves, perhaps using side payments from one to another, to see if they can all accept the computed highest-quality outcome. If they cannot all accept that negotiated solution, the process moves to the second-best outcome, and negotiation begins again. Hanging over the players, at every level in which negotiations are proceeding, is the threat of a worse negotiating position to come (for some of them) if they do not reach agreement there.

One-shot situations of the warfare type, in which the combatants have no interest in fairness and are trying their best to take *unfair* advantage of each other, do not fit into the framework of what I call game theory in this paper. Matrices to show the outcomes with various decisions made by the combatants can be constructed as guides to decision making, but the decisions will be guided by no rules--only by such things as gaining surprise, minimizing risk, and "fighting dirty". Cut-throat marketing and politics are often more like warfare than games.

Conclusion

Contest analysis started as a method I developed for evaluating military engagements between forces having many types of weapons on many types of vehicles. It worked so well there that it seemed like a

natural tool for other situations in which entities of various sorts are interacting in some quantifiable way. The athletic league was an obvious carryover from military combat and contest analysis works beautifully there. Next I tried situations where the interactions are between entities which are fundamentally different from each other--voters and candidates, for example. This also works beautifully, but people are not yet prepared to accept the idea of an indirect connection between competitors (election candidates interacting through voters). Here in game theory, we have another case in which the solution depends upon indirect interactions. The direct interactions are between strategies and outcomes; the strategies (and hence the players) interact with each other only indirectly through the outcomes.

There is nothing in common between this new quality-maximization method and traditional methods. At the risk of sounding grandiose, I have to characterize quality maximization as a new paradigm, replacing everything that went before in game theory.

Straffin makes the point that the Prisoner's Dilemma says something about the way human decisions often lead to what is called "the tragedy of the commons"--the idea that disaster can result when everybody acts in his own best interest. (When everybody can maximize his profit by grazing his animals on the public pasture instead of maintaining his own pasture, it eventually leads to the destruction of the "commons" and a worse situation for everybody.) Maybe Adam Smith was wrong, and maybe some experience with looking at game theory from the standpoint of fairness of outcomes will lead to better decisions. Straffin mentions the work of an author named Garrett Hardin, who suggests "mutual coercion, mutually agreed upon" as a way around the present overuse of the world's common resources. This coercion might be more acceptable if it is really an order to follow a set of rules coming from good game theory.

Acknowledgment

I am indebted to Prof. Philip D. Straffin for publishing *Game Theory and Strategy*. This book was invaluable to me in developing a small understanding of traditional approaches to game theory.

Alan E. Johnsrud
(Last revision: 10 December '11)

Go to Home page: [Home](#) Go to Topics page: [Topics](#)
A general resource for materials on game theory is at [resource](#).