Game Theory:
Modeling Decisions in Conflict

Sean H. Whalen
shwhalen@ucdavis.edu

March 20, 2009

1 Introduction

Human behavior often involves systems of conflict. Examples range from simple card games, to employees skipping work, to generals making decisions in battle. In each, two or more people are involved in making decisions whose outcomes conflict with each others' interests – the employee who feigns illness to enjoy a day off, versus the employer who desires productivity and has a limited tolerance for unexpected illness.

Drawing from social sciences and economics, Borel and Von Neumann[1] developed a mathematical model for decision making in these systems. In this paper we will examine the basics of this model, termed *game theory*, and how it explains two common systems of conflict: zero-sum and non-zero-sum games.

To understand the motivation behind game theory, consider this adage regarding social effort: "the purpose . . . is the greatest possible good for the greatest possible number" [3]. It should be clear that what is good for one is often bad for another, especially in large populations. There is no single maximization function for "goodness", but rather many individual and conflicting functions. Game theory defines a method for individuals to obtain a desired outcome when intelligent opponents seek the same for themselves.

---

[1] Borel's contributions date to the early 1920's, while Von Neumann independently developed the theory in the 1930's and 40's. Von Neumann and Morgenstern published the first formalization of game theory in 1944 [2].

First, we will discuss the formal foundations of game theory. They are not necessary for understanding games defined by the theory, so inclined readers may skip to Section 5.

2 Utility Theory

Decisions can be made by individuals or groups, under certainty, risk, or uncertainty [1]. A decision made in certainty always has a known outcome to the decision maker. A risky decision is one where known outcomes have known probabilities such as betting on a coin toss. An uncertain decision involves potentially unknown outcomes and definitely unknown probabilities. The distinction between individual and group is somewhat artificial, as a group can behave as a collective individual. As this is not always the case, such a distinction gives the model flexibility.

When making decisions, some quantity is usually maximized (such as profit) or minimized (such as cost). Under certainty, this reduces to a system of constraints, solvable by linear programming. In fact, every linear programming problem has an equivalent two person zero-sum gain (see Section 5) and vice versa [1]. More interesting is what happens under uncertainty or risk. With risk, we know the probabilities $p_1 \dots p_n$ and can first approach minimization or maximization using the expected value:

$$E[x] = \sum_{i=1}^{n} p_i x_i$$

This essentially gives a weighted sum of values produced by each outcome $x_1 \dots x_n$. However, there is a problem with using the unmodified expected value. Consider a wager on a coin flip, where the winner is paid a million dollars by the loser. This may be a reasonable bet for a multi-millionaire, but could bankrupt anyone else.

There is a non-linearity in an individual's preference for a wager's outcome. The transition from one million dollars to zero is more dramatic than one hundred million to ninety-nine million. Alternatively, a million dollars means much more to a penniless individual than to a multi-millionaire. We need a function which transforms the raw value of an outcome, such as monetary gain, into something representing an individual's preference.

This preference is called the *utility*, and a function transforming an outcome to a preference is a *utility function.* This function has certain formal consistency requirements proposed by Von

Neumann and Morgenstern [2]:

1. Any two outcomes must be comparable in terms of preference.
2. Preferences are transitive (if $A > B$ and $B > C$, then $A > C$).
3. A wager whose outcome is another wager (said to be *compound*) can be probabilistically separated into individual wagers.
4. If a person is indifferent towards two wagers, they are interchangeable in a compound wager.
5. If two wagers share identical preferred outcomes, the wager with a higher probability of occuring is preferred.
6. If there exist outcomes A, B, C where $A > B$ and $B > C$, a wager exists involving A and C where a player is indifferent to B.

This formalism simply says that to determine preference, the value of an outcome needs to be scaled by its' net worth to an individual (see Figure 1). Any pair of outcomes can be numerically compared in terms of utility. In fact, there exists a linear transformation $au + b$ between outcomes, where u is the first outcome's utility and a, b are constants.

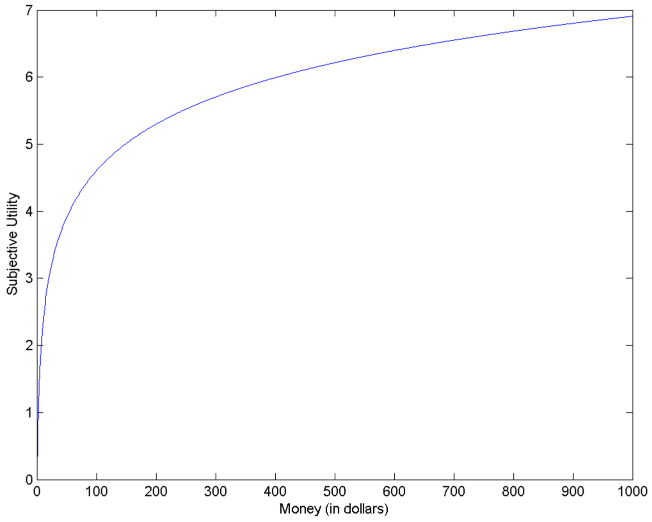

Figure 1: Money (in dollars) vs. subjective utility. The more money a person has, the larger an increase must be to have the same utility.

It is important to keep in mind that an outcome has a larger utility because it is preferred, and not vice versa. Also, because preferences are relative to an individual, they cannot be meaningfully compared between people.

3 Game Trees

Reasoning about games of any complexity is difficult. A *game tree* is a visual abstraction which relates all elements of a game, easing the process of analysis. An example game tree is given in Figure 2.

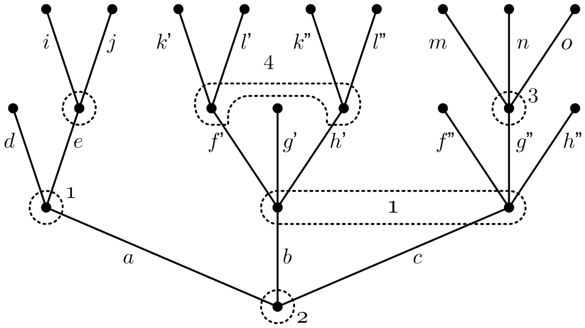

Figure 2: An example game tree from [1] (thanks to Sophie Engle)

The game progresses from the root node to a leaf node along a series of branches. A player selects one branch at each node. The final outcome is the value of a leaf node.

Our example tree represents only a subset of all possible moves. A complete tree would be prohibitively large for most games, due to combinatorial explosion. There are 400 chess games ending in two moves, while there are 197,742 ending in five[2].

Here, the game begins with player 2 selecting one of three moves. The probability of each move is represented by variables a, b, c. If the player selects the left branch, the next move is determined by a chance operation such as drawing a random card. The player cannot distinguish between the other two branches, due to some limitation such as incomplete information.

Game trees are simply another formalism, enabling a precise analysis of game progression.

---

[2] http://mathworld.wolfram.com/Chess.html

4 Rules of the Game

We will now formalize game trees and assumptions made about players. These are the *rules of the game* [1], which require the following:

1. A game tree of finite length, where nodes represent moves and branches represent potential outcomes (see Section 3).
2. A label for each node, representing which player (or chance itself) makes the move.
3. A probability for each move left to chance.
4. A division of moves into subsets, called *information sets*, where each set contains moves that cannot be distinguished by the player due to lack of information.
5. A labeling of each node in the game tree corresponding to the destination of a move in an information set.
6. A final outcome value for each endpoint in the tree.
7. A linear utility function for each player called the *payoff function.* This function defines the preference of the player over the set of endpoints.
8. Each player has full knowledge of the rules and payoff functions of all players.
9. Each player attempts to obtain the outcome with the highest utility, i.e. maximize their payoff function.

Rule 8 is unrealistic in practice, but gives the theory a sharp bound for modeling behavior in the most strict scenario.

5 Zero-Sum Games

The *zero-sum noncooperative game* is a simple system where gain for one player is balanced equally by loss to another. The sum of all gains and losses is zero. More formally:

$$\begin{aligned} \sum \text{gain} &= -\sum \text{loss} \\ \sum \text{gain} + \sum \text{loss} &= 0 \end{aligned}$$

Cooperation among players complicates the theory, and is outside the scope of this paper.

5.1 Uncertainty

If we view the world as completely determinate we can define probability theory as Glimcher does, "the tool we use to describe those portions of the environment about which we have incomplete knowledge." He continues, "By gathering more information we can reduce the uncertainty we face, and thus reduce our reliance on probabilistic models of the world, but we accomplish this at a cost in time and energy." [3]

Consider a game where the goal is prediction of a random machine-flipped coin toss. You could collect data about the machine and the coin, such as the center of gravity, and refine these measurements over repeated trials. By doing so, you could construct a probability distribution which increasingly reduces your uncertainty. As you approached certainty, you would always predict the same outcome when the same priors are observed. This is called a *pure strategy.*

5.2 Matching Pennies

The game of matching pennies illustrates a different kind of uncertainty, which is key to understanding the dynamics of a game system. Two players place pennies heads- or tails-up on a table. If both pennies land with the same face up, player 1 wins. Otherwise, player 2 wins.

As discussed in Section 5.1, player 1 collects data on the behavior of player 2. If she detects a pattern, she gains an advantage in the game. However, unlike a simple coin flipping machine, player 2 can disrupt player 1's pattern detection by introducing randomness in her behavior. From player 2's perspective, no reduction in uncertainty should be given to player 1. A formulation of the balance between strategy and randomness is a key contribution of game theory, and will be developed with more complex examples.

5.3 Holmes and Moriarty

A classic zero-sum game from [2] is presented in Table 1. This presentation style is called *strategic form*. The rows describe possible moves by Sherlock Holmes, and the columns describe possible moves by his archenemy Moriarty. Each cell in the matrix describes the payoff to Holmes and Moriarty for their choices.

The scenario has Holmes traveling via train from London to Dover, then to the continent to escape his pursuer Moriarty. Once he boards the train, he sees Moriarty outside. Holmes assumes he was spotted, and that Moriarty could overtake his train via an alternate route. Holmes then must decide to either continue to Dover or stop at the only intermediary, Canterbury. Moriarty anticipates this and must also make a decision, independent of Holmes. If the two meet at the same location, Moriarty will kill Holmes. If Holmes reaches Dover, he will escape to the continent.

	1: Moriarty goes to Dover to intercept Holmes at ferry	2: Moriarty goes to Canterbury to intercept Holmes on train
1: Holmes gets off train in Dover to escape to the continent	Holmes: -100 Moriarty: 100	Holmes: 50 Moriarty: -50
2: Holmes gets off train in Canterbury to escape Moriarty	Holmes: 0 Moriarty: 0	Holmes: -100 Moriarty: 100

Table 1: Sherlock Holmes game, from [2]

Table entries $(1, 1)$ and $(2, 2)$ given by (row, column) represent Moriarty catching Holmes. Accordingly, Moriarty receives a large gain and Holmes a large loss. Entry $(1, 2)$ represents Holmes escaping Moriarty, with a smaller gain to Holmes and loss to Moriarty (since escape is not as final as death). Entry $(2, 1)$ represents a tie, with Holmes escaping Moriarty at the intermediary, but failing to escape to the continent.

The gain and loss of each outcome sum to zero. However, Holmes is likely to lose more than he gains if he uses a pure strategy. Instead, a probabilistic strategy should be used by each player to avoid patterns of behavior which reduce uncertainty for the opponent. Using equations from [2], the optimal strategies are:

Moriarty goes to Dover	60%
Holmes goes to Dover	40%
Moriarty goes to Canterbury	40%
Holmes goes to Canterbury	60%

These probabilities are determined by the gain and loss values in Table 1. Such probabilistic behavior is called a *mixed strategy*. A mixed strategy can be thought of as an assignment of probability to each possible pure strategy. Furthermore, every pure strategy is a mixed strategy, with one choice at 100% and all others at 0%.

5.4 Equilibrium

An optimal strategy in the context of game theory allows a player to maximize gains and minimize loses. The values of an optimal mixed or pure strategy are *equilibrium points*, said to form an *equilibrium strategy*.

Following an equilibrium strategy means neither Holmes nor Moriarty should vary their behavior from equilibrium points unless the other changes strategy. By straying from equilibrium points a player may obtain larger gains, but does so at the risk of larger losses. If a player takes this risk, it is then better for the opponent to also adopt a non-equilibrium strategy [3].

Multiple equilibrium points can co-exist, and have equal utility in zero-sum games. Not all games have an equilibrium strategy [1].

6 Non-Zero-Sum Games

In *non-zero-sum noncooperative games*, player gains and losses are not in balance. More formally:

$$\begin{aligned} \sum \text{gain} &\neq -\sum \text{loss} \\ \sum \text{gain} + \sum \text{loss} &\neq 0 \end{aligned}$$

6.1 The Prisoner's Dilemma

Two suspects are arrested and interrogated in separate rooms. Each is given the option of confession. If neither confess, they will serve some time on a trumped-up charge. If both confess, they will receive a slightly lesser sentence. If one confesses and the other does not, the confessor will serve a light sentence and the other a full sentence. A strategic form is given in Table 2. Notice the asymmetry in gains and losses.

	1: Prisoner B holds out	2: Prisoner B confesses
1: Prisoner A holds out	Prisoner A: 1 year Prisoner B: 1 year	Prisoner A: 10 years Prisoner B: 3 months
2: Prisoner A confesses	Prisoner A: 3 months Prisoner B: 10 years	Prisoner A: 8 years Prisoner B: 8 years

Table 2: The prisoner's dilemma, from [1]

Neither suspect can know what the other will do. The utility of a mutual confession is less than the utility of a single confession, providing the temptation for betrayal.

A "rational" suspect will not make a decision dependent on the other. Regardless of Prisoner B's decision, Prisoner A will be better off confessing. However, if one suspect thinks this way, so will the other. They will both confess and serve a nearly full sentence.

The equilibrium strategy for both suspects is confession. But if they both behave "irrationally", they serve the least overall time. Straying from the equilibrium chances larger utility and risks larger loss.

The Prisoner's Dilemma is a classic example of a non-zero-sum game. It can be played online to gain greater intuition [3], but for now we will move on to a different example.

---

[3] http://serendip.brynmawr.edu/playground/pd.html

6.2 Chicken

Table 3 contains the strategic form for the game of chicken. Two players race towards each other in vehicles, the loser being the first to swerve and avoid collision. If Jones swerves before Smith, he loses 10 points and Smith gains 50. Vice versa if Smith swerves first. If both swerve, there is a negligible gain for each since they live to play again. If neither swerve, there is a severe penalty for both as they will collide and risk death.

	1: Smith continues	2: Smith swerves
1: Jones continues	Jones: -100 Smith: -100	Jones: 50 Smith: -10
2: Jones swerves	Jones: -10 Smith: 50	Jones: 1 Smith: 1

Table 3: Game of chicken, from [3]

To maximize their score, the players develop a mixed equilibrium strategy. This takes place over repeated trials, ignoring the possibility of death. First, we note that:

$$\begin{aligned} P(\text{Smith continues}) + P(\text{Smith swerves}) &= 1 \\ P(\text{Smith continues}) &= 1 - P(\text{Smith swerves}) \\ P(\text{Smith swerves}) &= 1 - P(\text{Smith continues}) \end{aligned}$$

Where $P(x)$ indicates the probability of event x occurring. Furthermore, using the formula for expected utility we can say:

$$\begin{aligned} \text{If Smith continues, Gain to Jones for swerving} &= P(\text{Smith continues}) * -10 \\ \text{If Smith swerves, Gain to Jones for swerving} &= P(\text{Smith swerves}) * 1 \\ \text{Gain to Jones for swerving} &= P(\text{Smith continues}) * -10 + \\ & \quad P(\text{Smith swerves}) * 1 \\ \text{Gain to Jones for swerving} &= \big(1 - P(\text{Smith swerves})\big) * -10 + \\ & \quad P(\text{Smith swerves}) * 1 \end{aligned}$$

If Jones is to behave according to an optimal mixed strategy, he must find an equilibrium point where swerving and continuing have equal utility. Continuing our equations:

$$\begin{aligned} \text{Gain to Jones for continuing} &= \big(1 - P(\text{Smith swerves})\big) * -100 + \\ & \quad P(\text{Smith swerves}) * 50 \end{aligned}$$

We can now find the equilibrium point by setting Jones' gain for swerving equal to his gain for continuing:

$$\begin{aligned}\big(1 - P(\text{Smith swerves})\big) * -100 + P(\text{Smith swerves}) * 50 \quad &= \\ \big(1 - P(\text{Smith swerves})\big) * -10 + P(\text{Smith swerves}) * 1 &\end{aligned}$$

The solution is $P(\text{Smith swerves}) = 0.647$, meaning if Smith swerves 64.7% of the time, Jones assigns equal utility to continuing and swerving. Each choice is equally good or bad, which has reduced his uncertainty. This remains optimal behavior as long as Smith is following an equilibrium strategy.

7 Summary and Conclusion

Game theory mathematically describes optimal behavior for maximizing gains and minimizing losses when intelligent players are in conflict. It does this through formal analysis of different game types. The rules of the game establish a game tree, assumptions about players, and the players' relative preferences for outcomes.

Zero-sum and non-zero-sum games describe systems where player gains are matched by equal and unequal loss, respectively. We discussed several examples of each. Variations we did not discuss allow for elements such as cooperation among players and group decision making.

A player can decrease uncertainty about opponents' behavior by collecting data and detecting patterns. An opponent can counter by introducing probabilistic behavior to increase uncertainty. Mixed strategies find points of equilibrium, probabilities which suggest optimal behavior against equilibrium-abiding opponents. A player can stray from equilibrium to risk larger gain and larger loss. A mixed strategy with a single behavior is a pure strategy.

Game theory was originally a tool for economists, but has found applications in areas such as neuroscience, psychology, and sociology. This paper is merely an introduction, and a wealth of knowledge exists on the subject for the interested reader. Von Neumann and Morgenstern's classic text [2] is geared towards the mathematically inclined, while Luce and Raiffa's text [1] is more abstract. Paul Glimcher's text [3] provides excellent motivation and discussion of the theory, though is much broader in focus.

References

[1] R.D. Luce and H. Raiffa, *Games and Decisions: Introduction and Critical Survey*, Dover Publications, 1957

[2] O. Morgenstern and J. Von Neumann, *Theory of Games and Economic Behavior*, Princeton University Press, 1944

[3] P.W. Glimcher, *Decisions, Uncertainty, and the Brain: The Science of Neuroeconomics*, Bradford Books, Feb. 2003