

Game theory for beginners: On lazy students, poker and group dynamics

Michael Muskulus*[†]

Wednesday, 14 November 2007

Abstract

Game theory is a branch of mathematics with interdisciplinary applications in areas as diverse as our daily life, evolutionary biology and our economy. Not surprisingly, one of the Nobel prizes of this year has been awarded due to research on game theoretic questions (Hurwicz, Maskin & Myerson). In this short lecture I will discuss some basic notions and models of game theory, whose knowledge might not only be of use to students, but who, in my opinion, should be known by every serious scientist.

*Leiden University, Niels Bohrweg 1, 2333 CA Leiden, The Netherlands
Email address: muskulus@math.leidenuniv.nl (M. Muskulus)
Web: www.math.leidenuniv.nl/~muskulus/

[†]These are notes of an invited talk given in front of student society “De Leidse Flesch” in Leiden, The Netherlands.

1 Introduction

Ladies and Gentlemen, thank you for giving me the opportunity to talk in the relaxed setting of this “lunch lecture”. Usually I am an adherent to the free presentation, but for this special occasion I am experimenting with the very traditional way of the well prepared lecture that asks to be read off from paper. One reason for this decision is the time constraint of 30 minutes, which is extremely short when we want to talk about mathematics. The german mathematician *Klaus Jänich*, author of popular textbooks and famous for his informal writing style, once remarked (if I remember that correctly) something in the spirit of: that one of the more important mistakes students make when faced with real mathematics for the first time is a profound confusion about the time needed to do things properly. According to him, reading a page in a well-written mathematical textbook has to be measured in hours, not in minutes. To which I would like to add the words of *Paul Halmos*, from his autobiography, which goes by the title “I want to be a mathematician”, that for each minute of a professional presentation given one needs one hour to prepare it.

But anyway, why is mathematics so difficult? This is, above all, a human problem. Epistemologically — when considering it from the viewpoint of how it deals with knowledge — mathematics is almost¹ trivial. All knowledge in mathematics is deduced by logical proof, so the reason that mathematics is problematic for many people, is that humans are epistemologically incomplete subjects. By this I mean, that we often don’t know what we *know*, what is already given to us, in the form of our beliefs and assumptions, but rather often need the help of a friend playing Socrates, or the devil’s advocate, to understand the implications of our thoughts. This is an important problem — and one students should be familiar with. After all, the process of learning is, for a great deal, being *surprised* by what you actually *can* do, even with limited tools and under time constraints.

Alas, let me now apply the mathematical method to this discussion and henceforth assume that people not only know where their thoughts lead to, but also that they know what their *actions* entail; at least in so far, as to how much the consequences apply to them. It would then be *rational* to choose your actions according to the desirability of the expected outcomes. A very practical problem that presents itself in this theory of *rational choice* then is the question how to choose between alternative actions. Obviously, we need to order the possible outcomes according to their desirability, and this order

¹ This is *Kurt Gödel’s* catch, of course: A logical system of such complexity as mathematics will either be inconsistent, or there will necessarily exist truths that cannot be proved by deduction [11].

should be total and transitive: If I prefer A to B, and B to C, then it would be problematic if I would also prefer C to A. In the language of *Stanley Smith Stevens* theory of scales of measurement [13], I assume that my preferences are *ordinal* measurements.

In fact, as a mathematician I want to assume that they are *interval* measurements, such that I can assign numerical values to my preferences. The unit in which these are expressed is the so-called *util*, since what is measured is the utility of the results to me, that is, to speak with the words of *Jeremy Bentham* [4], the pleasure that I derive from them. Note that pleasure should be understood ethically here, that is, in its use for the “good life” that I want to lead [2]. In particular, it does not endorse short-term joy or selfish behaviour, but encompasses a rather complete evaluation of the impact on my life’s situation with respect to all its aesthetic, emotional, social and ethical values. In economics, the function that tells us the expected utility that arises from a particular outcome is, nevertheless, called a payoff function, but we must not forget that although money is important as an incentive for people, utility is usually much more.

The method to find the actual numerical value of our payoff function in a given situation is the same as the one proposed by our friends the Bayesians, who consider probabilities as subjective estimates [1], when they want to know how probable we think something is — and as it happens in some peculiar television shows: assume we prefer A to B, but we are offered B plus some amount of additional utility X, how much of this extra utility do we need to be *indifferent* between A on the one hand, and B plus X on the other hand? How much persuasion would you need to accept an “indecent proposal”? Or how much money?

Unfortunately, human nature can lead to even more serious problems, namely, the sin of being inconsistent. The following example is taken from Osborne [9]. Assume you partake in a meta-lottery that offers you the choice between two outcomes.

- (1) You receive two million Euro for sure.
- (2) You receive:
 - ten million Euro with probability 0.10,
 - two million Euro with probability 0.89,
 - nothing with probability 0.01.

Which of the two lotteries (1) or (2) do you prefer? Then consider the following two choices:

(3) You receive

- two million Euro with probability 0.11,
- nothing with probability 0.89.

(4) You receive:

- ten million Euro with probability 0.1,
- nothing with probability 0.9.

Which do you prefer now? A significant number of experimental subjects say that they prefer lottery (1) to lottery (2), and lottery (4) to lottery (3), but these preferences cannot be represented by an expected payoff function. To be sure, let us prove this: Assume there exists a payoff function $u : \mathbb{R} \rightarrow \mathbb{R}$ such that

$$u(2) > 0.1u(10) + 0.89u(2) + 0.01u(0).$$

Subtracting $0.89u(2)$ and adding $0.89u(0)$ to each side we obtain

$$0.11u(2) + 0.89u(0) > 0.1u(10) + 0.9u(0),$$

which is a contradiction, since this means we should prefer lottery (3) to lottery (4).

A further problem is that we are not alone in this world. I don't want to consider moral questions here, but the existence of at least another person who might have different ideas about what kind of outcome she would prefer in a given situation, means that we cannot always have the best result for us. Now game theory is nothing else than the interdisciplinary effort to make sense of such situations. Or to say it in the words of *George Spencer-Brown*: "Only two can play this game" [12].

2 Strategic games

In its mathematical formulation, a strategic game consists of a set of players, a set of actions for each of them, and a payoff function for each player that depends on the strategic choices of all players involved. A two person game with finite actions therefore, is equivalent to two matrices, one for each player. Note that games like chess, go or poker usually consist of many moves that depend on the history of moves played so far by all players. But these games, in their so-called *extensive form*, can also be written in a strategic form, if we make the payoff matrices large enough to cover all possible histories of moves.

Let us now consider a famous game which might be called **Working on a Joint Project**: you and a fellow student work together for a course, and each of you can choose to either work very hard, or to take it easy and goof off. You both want to pass the course, of course, but you both don't like working very much. Assume the following payoff matrix, where the first number in each entry is the payoff for you, and the second one is the payoff for your partner:

Your choice:	Your partner's choice:	
	Work hard	Goof off
Work hard	2, 2	0, 3
Goof off	3, 0	1, 1

These numbers assume that you prefer that you both work hard to the situation in which you both goof off. But the best outcome would be if you could take it easy, and your partner would do the work. The worst possible outcome would be if you were working hard, and your partner would not, since you don't want to be exploited.

Assume that you both are going to work (or relax) on your own today, and then will meet tomorrow to combine your results. What would you do? And moreover, can you decide this rationally, such that you can explain and recommend your behaviour to other students in the same situation?

The mathematical genius *John von Neumann* advocated the use of the so-called *maximin* solution, which is to choose that strategy that guarantees you the maximum possible outcome for you, even when your opponent could have guessed your choice. In this pessimistic world-view, your minimal payoff would be

Your choice:	Your minimum possible payoff:	
Work hard	0	(Goof off)
Goof off	1	(Goof off)

so the maximin strategy would be to goof off, which would, at least, guarantee you one util. Might this be the reason why it is so difficult to work productively together?

John Nash, the tragic hero of Sylvia Nasar's book "A beautiful mind", was awarded the 1994 Nobel price for a different approach which tries to be more optimistic. Assume we know what our partner is going to do. Then our *best reply* strategy would be the following:

Your partner's choice:	Your best reply payoff:	
Work hard	3	(Goof off)
Goof off	1	(Goof off)

and likewise for our partner. A *Nash equilibrium* is then a strategic choice for each player such that every player's action is a best response to the other players' actions, and we can interpret this as a *stable social norm*, to which experienced players should adhere. Or as *Ken Binmore* has stressed, a Nash equilibrium allows us to break out of the circular reasoning of the kind "I think that you think that I think that you..." and find a solution of the underlying equation. Unfortunately, in the case of the **Joint Project** game, only the action (**Goof off, Goof off**) is a Nash equilibrium, and this is the same as the pessimistic maximin solution. In fact, this game is a so-called game of *pure conflict*, also known as the **Prisoner's dilemma** or the **Arms race**, and there is nothing we can do to avoid the sub-optimal solution of lazy teamwork.

A related game is the so-called **Battle of the Sexes**, which is modelled by the following payoffs:

Your choice:	Your partner's choice:		
		Football Match	Classical Concert
Football Match		2, 1	0, 0
Classical Concert		0, 0	1, 2

So we assume that, one a given night, you can either go to see a football match, or a classical concert. Now you prefer the football match, but your partner prefers the concert. But the worst possible outcome would be to have a difficult argument and not go together, or at all. So what should you do?

The maximin strategy does not help you at all, since it is zero, no matter what action you choose. But if you knew that your partner will choose the classical concert, your best reply is to choose the same. And if she chooses the football match, your best reply is to do the same. So here we have two Nash equilibria, with payoffs (2, 1) and (1, 2), respectively. They correspond to two stable social norms, namely: **Ladies first**, in which your partner always gets what you know he wants, and **Gentlemen first**, in which you always get what your partner knows that you want. Both of these actions are unfair, but this cannot be helped in such a game.

But one problem remains: which of the two equilibria is going to be played? When faced with this issue, the American economist *Tom Schelling* proposed that we pay attention to so-called *focal points*; work for which he received the 2005 Nobel price. The idea is simple and fascinating at the same time. I will give two examples.

1. Assume you and a friend want to meet tomorrow at 12 o'clock in Amsterdam, but you both have never talked about exactly where. Now it

is too late, you cannot reach him anymore, and also you are impossible to reach, for some reason. Where would you go to, tomorrow at 12 o'clock?

2. Assume you and a friend partake in a quiz show on television. You are told that you both have to choose either **Heads** or **Tails**, independently of each other, and if you happen to choose the same, you both will get 100 Euro, and otherwise nothing. What would you choose?

I am not saying that this idea will solve the **Battle of the Sexes**, but it should give you some hints on possible solutions that might work in many practical situations.

3 Mixed strategies

Let us now, very briefly, consider a game called **Matching Pennies**, which is very simple. You and your opponent both choose either **Heads** or **Tails**, and if you both say the same, then you will win one Euro from your opponent, otherwise he will win one from you:

Your choice: Your opponent's choice:

	Head	Tail
Head	1, -1	-1, 1
Tail	-1, 1	1, -1

This game is quite interesting. For a start, it does not have a (pure) Nash equilibrium. Moreover, if you play this game more than once, and against the same opponent, then it is easy to see that the only way not to lose large amounts of money, is to be *unpredictable*. So you should consider playing a probabilistic strategy — or at least a deterministic, chaotic strategy, which for all practical purposes can be assumed to be the same.

As a related problem, imagine that your opponent can masterfully read off from your body language what you are going to do — or even that he is using a lie detector. Then the rational thing to do would be not to know yourself what you are going to do: so you throw a dice, or even better, you use a quantum mechanical random generator. Mathematically, this is modelled by a *mixed strategy*, in which we choose our action by a probability law. For example, the strategy $(1/2, 1/2)$ means that we choose **Heads** 50 percent of the time, and **Tails** the rest of the time.

Assume we use a mixed strategy $(p, 1 - p)$, and our opponent uses a mixed strategy $(q, 1 - q)$, and we play **Heads** and **Tails** very often. Then the probabilities of our common choices are

Your choice: Your opponent's choice:

	Head	Tail
Head	pq	$p(1 - q)$
Tail	$(1 - p)q$	$(1 - p)(1 - q)$

and our expected payoff E is²:

$$\begin{aligned} E &= pq u_1(1, 1) + p(1 - q) u_1(1, 2) + (1 - p)q u_1(2, 1) + (1 - p)(1 - q) u_1(2, 2) \\ &= p(qu_1(1, 1) + (1 - q)u_1(1, 2)) + (1 - p)(qu_1(2, 1) + (1 - q)u_1(2, 2)). \end{aligned}$$

Filling in the payoffs for **Matching Pennies**, we get

$$\begin{aligned} E &= p(q - (1 - q)) + (1 - p)(-q + (1 - q)) \\ &= p(2q - 1) - (1 - p)(2q - 1) \\ &= (2p - 1)(2q - 1). \end{aligned}$$

Since I want to maximize this, I will choose the value of $p = 1/2$. Likewise, my opponent, which gets the negative of my payoff, wants to minimize this, so he will also choose the value of $q = 1/2$. In fact, the mixed strategy characterized by $(p, q) = (1/2, 1/2)$ is a *mixed Nash equilibrium*, since we both cannot do better than playing this strategy when our opponent does adhere to it.

Returning to the **Battle of the Sexes**, its payoff for us, the first player, is

$$E_1 = 2pq + (1 - p)(1 - q) = (1 - p - q) + 3pq,$$

and for our partner

$$E_2 = pq + 2(1 - p)(1 - q) = 2(1 - p - q) + 3pq.$$

The maximum of E_1 , considered as a function of p , is achieved for $q = 1/3$. Likewise, the maximum of E_2 , considered as a function of q , is achieved for $p = 2/3$. As before, the pair $(p, q) = (2/3, 1/3)$ is a mixed Nash equilibrium, and the expected payoff of this equilibrium is $E_1 = 2/3 = E_2$, so this is a *fair* solution to the battle of the sexes. Unfortunately, it is also inefficient, since $2/3$ is even less than the payoff of 1 we would get if we would, for example, play the **Ladies first** pure equilibrium strategy.

² Note that the first term in brackets is the expected payoff of the pure strategy of always saying **Heads**, and the second term likewise of the pure strategy of always saying **Tails**.

The "real solution", of course, is to use communication. But still: why should you trust on an agreement, if your partner can really profit if he eventually does something different than what he promised to do? What is really necessary, thus, is a *mechanism* that guarantees that both players do what they agreed on. In game theoretic parlance, this would allow the player to achieve a more efficient *correlated* equilibrium. Remarkably easy as this idea sounds, *Robert Aumann* was awarded the joint 2005 Nobel price, among other things, for working out the technical details of this.

4 Group dynamics

Consider the game variously called **Reporting a Crime** or **The Good Samaritan** [5]. It is played by n players who all observe a crime. Each person involved would like the police to be informed, but everybody prefers that *someone else* does call them. So each player can choose between the two actions **Call** or **Nothing**, and the payoff is 0 if nobody calls, a certain value v if at least one person else calls, and $v - c$ if oneself calls (and maybe somebody else, too). This game has no pure Nash equilibrium, since if everyone calls, it would be better for you not to do so. And if nobody calls, it would be better to call yourself.

However, there exists a mixed Nash equilibrium, which can be found by a symmetry argument: In any such equilibrium, each person's expected payoff when calling is equal to her expected payoff when not calling³. This means that we require $E[\text{Call}] = E[\text{Nothing}]$, such that

$$\begin{aligned} v - c &= 0 \cdot P(\text{no one calls}) + v \cdot P(\text{at least one other person calls}) \\ &= v(1 - P(\text{no one else calls})), \end{aligned}$$

or simply $c/v = P(\text{no one else calls})$.

The probability that no one else calls is the probability that every one of the $n - 1$ other players does not call. Denoting the probability with which each person calls by p , then $c/v = (1 - p)^{n-1}$, and finally

$$p = 1 - (c/v)^{1/(n-1)}.$$

As n increases, the value of p decreases. For example, if $v = 10$ and $c = 1$, then with only two players each should help with probability $p = 9/10$. With

³ The reason is, of course, that if one were not indifferent between the two *pure strategies* of **calling** or doing **nothing**, than one would not play the mixed strategy, but the more preferred pure strategy — remember that mathematical expectation is linear! But then the mixed strategy would not be an equilibrium.

three players, each should help with probability $p \approx 0.68$, with four players, $p \approx 0.54$, etc.

What about the probability P_0 that nobody calls the police? This is, of course,

$$P_0 = (1 - p)^n = (c/v)^{n/(n-1)} = (c/v)^{1-1/n}$$

and this also decreases with increasing n , asymptotically to the value c/v . So in this example, if a large number of people is involved, then in 1 case out of 10 the police is *not called* — the consequences of which I leave to your imagination to consider.

5 Playing Poker

Consider the popular game of Texas Holdem poker, in which players are dealt two *pocket cards* each, and five common cards are laid out, in three stages, on the table, for everyone to see. Let us consider a special game situation here⁴, adapted from Edward Packel’s book “The Mathematics of Games and Gambling” [10]: You are only facing one opponent, the last round of betting is about to begin, and we assume the following three variables:

- p = the probability of you having the better hand, intelligently estimated by your opponent
- k = the current pot size before the last round of betting
- r = the maximum allowable bet in the final round

We furthermore assume that your opponent will never raise you, only call or fold, and that you yourself know whether you have the better hand (in which case you would bet the additional amount r), or not (in which case you would either fold, or also bet the additional amount r in order to intimidate your opponent by bluffing). Your opponent then has to decide whether he calls your bet (by putting the same amount r into the pot), or whether he folds.

The game we are considering therefore has the following expected payoffs for pure strategies:

⁴ For more on game theory applied to poker, you could have a look at the homepage of the University of Alberta computer poker research group, <http://www.cs.ualberta.ca/games/poker/> and at the recent book by Chen & Ankenman [6].

	Your choice:		Your partner's choice:
		Always Call	Always Fold
Never Bluff	$\alpha = p(k+r) + 0$	$\beta = pk + 0$	
Always Bluff	$\gamma = p(k+r) + (1-p)(-r)$	$\delta = pk + (1-p)k$	

As before, we can find a mixed Nash equilibrium by the indifference argument⁵. Let x denote the probability that you are bluffing, and y denote the probability that your opponent is calling. For y at a mixed equilibrium, your expected payoff should be indifferent on your strategy:

$$E[\text{always bluff}] = \gamma y + \delta(1-y) = E[\text{never bluff}] = \alpha y + \beta(1-y).$$

Letting $D = \alpha - \beta + \delta - \gamma = (1-p)(k+r)$, this has solution

$$y = \frac{\delta - \beta}{D} = \frac{k}{k+r}.$$

Likewise, for x at a mixed equilibrium, your opponent's expected payoff, which is the negative of yours, should be indifferent on his strategy:

$$-E[\text{always call}] = \alpha(1-x) + \gamma x = -E[\text{never call}] = \beta(1-x) + \delta x,$$

which has solution

$$x = \frac{\alpha - \beta}{D} = \frac{pr}{(1-p)(k+r)}.$$

Finally, your expected payoff under this equilibrium is, after some algebra,

$$E = \alpha y + \beta(1-y) = \frac{\alpha\delta - \beta\gamma}{D} = \frac{pk(k+2r)}{k+r}.$$

Let us apply this to an example. Assume the pot has value $k = 100$, the maximum possible bet is $r = 100$, and you estimate $p = 1/5$. Then $D = 160$ and when you do not have the winning hand, you should bluff with probability $x = 1/8$, and your opponent will call with probability $y = 1/2$. Your expected payoff will be 30. If you bluff less — or more often — and your opponent suspects this, she can change her strategy and your expected payoff will be less than this. The main lesson to be learned from game theory

⁵ A technical point on the side: In order for this to apply, one need to make sure that so-called *dominated strategies* have been eliminated. This happens whenever $p < (k+r)/(k+2r)$, and only then does our simple analysis apply. In the other case, your opponent should never call, and you should always bluff [10].

when applied to poker is that you should bluff quite often — usually more often than people normally do!

On the practical side, how can one be sure to bluff with the right probability? People are notoriously bad at generating random numbers [3], and if you would throw the dice, your opponent is likely to assume that you are bluffing, unless you do this every time before you play! A practical way is to base your decision on some chance event that is displayed publicly, on the table in front of you. For example, consider the last card that has been added to the five common cards, which is not an eight. There are $52 - 4 = 48$ possible cards⁶ for this, so if this card is a heart of rank higher than eight, you would bluff; and otherwise, you would not. Then you would bluff with probability $6/48$ (that is, if the card would be one of the nine, ten, jack, queen, king or ace of hearts) as desired. Any other such scheme is also possible. The only thing you have to make sure of is that your opponent does not know how exactly you do this.

6 A final problem

Before I want to end with a last problem, let me just mention the many applications of game theory I could not talk about. These cover, among other things,

- auctions (à la eBay and beyond),
- voting (in all its forms),
- and evolutionary biology.

But now consider the **Surprise test paradox**, which is problem number seven in *Martin Cohen's* famous book of philosophical problems [7]: A teacher tells his class that they are to be given a test one day next week, but the day on which the test will be given will come as a surprise to them. Now Bob is a pupil who has the following idea: If the test were given on the last possible day, Friday, it would not be a surprise. So this can't be. But then assume the test were given Thursday. Since it cannot be given on Friday, if it is not given before Thursday, it will be no surprise on Thursday. So this also cannot be. Continuing in this way, technically this is called *backward*

⁶ In fact, this is only an approximation, since you are likely to have much more information about the cards that could appear — or rather, those that could not. For instance, your two pocket cards cannot appear on the table in front of you. So to correct this, one needs a scheme which is a little more elaborate — but based on the same idea.

induction, Bob reasons that the test cannot be given at all. Consequently, he does not learn. A few days pass, but on Wednesday the teacher announces that the class will now have to take the test. Bob says: “But you cannot do this!” The teacher says: “Why not?” Bob goes on: “Because the test needs to be a surprise, so you can only hold it when we are not expecting it!” Teacher: “Yes, but it seems to me, Bob, that you are *not* expecting it, and I *am* holding the test...”

Now, as *Ken Binmore* suggests [5], the main problem arises because Bob and the teacher think they play the same game, whereas in fact it is two different games. What are they?⁷ And what would you do?

Thank you!

⁷ There is also a nice analogy with the so-called *ultimatum minigame* [5], but unfortunately this would take us too far here.

References

- [1] G. D'Agostini, Teaching statistics in the physics curriculum: Unifying and clarifying role of subjective probability, *American Journal of Physics* **67** (1999), 1260–1268. Preprint: <http://arxiv.org/abs/physics/9908014>.
- [2] Aristotle, *The Nicomachean Ethics*, Penguin Books (2004).
- [3] S. A. Armstrong, *A Meta-Analysis of Randomness in Human Behavioral Research*, MSc Thesis, Department of Mathematics, Louisiana State University (2004). URL: <http://etd.lsu.edu/docs/available/etd-07092004-113059/>.
- [4] J. Bentham, *An Introduction to the Principles of Morals and Legislation*. Dover (2007), originally published 1789.
- [5] K. Binmore, *Game Theory: A Very Short Introduction*. Oxford University Press (2007).
- [6] W. Chen & J. Ankenman, *The Mathematics of Poker*. ConJelCo (2006).
- [7] M. Cohen, *101 Philosophy problems*. Routledge (2002).
- [8] P. R. Halmos, *I want to be a mathematician. An Automathography*. Springer (1985).
- [9] M. J. Osborne, *An Introduction to Game Theory*. Oxford University Press (2004).
- [10] E. Packel, *The Mathematics of Games and Gambling*. The Mathematical Association of America (2006).
- [11] P. Smith, *An Introduction to Gödel's Theorems*. Cambridge University Press (2007).
- [12] J. Keys (G. Spencer-Brown), *Only Two Can Play This Game*. Julian Press (1972).
- [13] S. S. Stevens, On the theory of scales of measurement. *Science* **103** (1946), 677–680.