On the Need for an Inductive Approach to Game Theoretic Questions

Ben Labe
University of Pittsburgh
bdl15@pitt.edu

December 14, 2009

Abstract

The following paper advocates an alternative approach to game theoretic problems. In order to make itself more scientific, Game Theory should be concerned with meeting the following three criteria: Adequacy, Efficacy, and Prescriptiveness. Sometimes, the greatest way to satisfy these three criteria is not to work from the universal to the existential, but rather from the existential to the general. A simple, hypothetical treatment of the Prisoner's Dilemma will be offered which follows the method proposed. This should have the dual effect of explaining some of the experimental data surrounding the Prisoner's Dilemma and providing the motivation for adopting the inductive method. This project will also require an interpretation of the normative status of terms like 'solution concept' and 'rationality'.

1 Introduction

The Prisoner's Dilemma is well-known because it illuminates two facts. The first is that game-theoretic solution concepts do not always coincide well with empirical data. The second is that many people are willing to forfeit personal gains for communal ones in games of strategic interaction. That these two facts cause problems for Game Theorists should come as no surprise. Most current solution concepts rest upon a particular assumption of rationality, which cannot account for the proportion of people who are not merely self-interested. However, a new branch of literature, which focuses on psychological games, has been trying to address that discrepancy.

While psychological game theory has already accomplished much, it has very far to go. Part of the purpose of this paper is to highlight some of the problems that currently face game theorists, and to state how I think they can be resolved. Ultimately, I believe that the answer to those problems rests in an inductive, psychological approach to games. This may be the only way to explain the Prisoner's Dilemma and other paradoxes like it. Ultimately, it rests in part on an understanding of how to properly interpret concepts like normativity and rationality from within and without the field.

Early on, it might be important to assure people that the goal of this essay is not to discredit the entire practice of the game theory that is built on the assumption of self-interestedness; nor is it to provide a solution to the Prisoner's Dilemma that advises players to cooperate. Game theory is perfectly capable of describing and predicting how people will act in a variety of situations. Insofar as it deals strictly with the assumption of self-interestedness, however, it is only able to account for the proportion of the human population that is merely self-interested. Since this proportion usually constitutes a majority, game theory has done much to advance our understanding of how people operate. I only wish to advise game theorists to be wary of how they interpret some of their results. Such an awareness is particularly important whenever one is offering prescriptions. To try to provide a solution concept that rejects the Nash Equilibrium strategy would be against the very spirit of this work. The goal here is to develop a way of viewing the Prisoner's Dilemma that allows us to predict and to explain the behavior of as many types of people as possible. This cannot be done if we ignore a large portion of the population. Moreover, any attempts at making prescriptions should come as secondary, and should always be treated hypothetically. Thus, the matters of primary importance to game theorists should be adequacy and efficacy.

This paper comes in two parts. In the first, I will justify an inductive approach to the problem of adequacy in the Prisoner's Dilemma. In the second, I will actually show what such an approach might look like. Insofar as it is possible, I have tried to restrict the sources substantiating my arguments to those within game theory itself. However, due to the nature of the first part of the essay, at times I will resort to arguments and concepts from within philosophy. Sometimes, you can only take down a city by attacking it from the outside.

2 The Strength of a Theory

1. Adequacy, Efficacy, and Prescriptiveness

Any good scientific theory seeks to satisfy the criteria of adequacy (the ability to describe how things *do* happen), efficacy (the ability to predict how things *will* happen), and prescriptiveness (the ability to specify how things *should* happen). In the natural sciences, in which the phenomena being studied are not human (*a fortiori*, 'rational'), the latter two criteria amount to the same thing. For example, a chemist who has sufficient knowledge of two chemicals might be justified in rejecting the results of an experiment by appealing to what *should* have happened when the two chemicals were mixed. An aberration of this sort will often be looked upon as a problem of experimenter error before it ever will be looked upon as a counterexample to the general rule. For the chemist, that is a problem of efficacy rather than prescriptiveness.

The economist experiences a special advantage over the natural scientist. Since the economist is dealing with human behavior, the distinction between what constitutes efficacy and what constitutes prescriptiveness becomes ever so clear. In economics, the notion of what ought to happen takes on normative force, to which economists can hold subjects obligated. In short, economists can *blame* subjects who do not follow their advice.

Such an advantage should not be treated lightly. It comes as a disadvantage, too. The problem for the economist is that the subjects of his experiments do not always act as mechanistically as do the subjects of the natural sciences; or else, if they do, there are too many factors at work to determine exactly how that mechanism might operate.

As such, the game theorist is stuck in the delicate position of having to describe both how players *do*, or *will*, behave, and how they *should* behave, both at the same time. Meanwhile, those may be separate matters entirely. Thus, the economist must make sure to distinguish between the type of 'should'-statements that invoke efficacy, and the type of 'should'-statements that invoke prescriptiveness. I will often do this by referring to the second type of statement as an 'ought to' rather than a 'should'.

The strength of Game Theory is most obviously reflected by its solution concepts. Ritzberger [11] has the following to say about their role (p. 189):

> Abstractly, a *solution concept* is a mapping from the set of games to outcomes or strategy combinations. It assigns to every game one or several "solutions". There are various ways to interpret such a mapping. It can be thought of as a rule which gives normative *recommendations* to rational players on how to behave in any given game. It may also be viewed as the game theorist's *predictions* of what might happen in the various games if those were played by rational players. In practice, it is often a way of arguing why certain observed phenomena exist.

Thus, in providing adequacy, the game theorist simultaneously provides explanation. We can now see the bridge that is created between adequacy and prescriptiveness. Not only should a game-theoretic model supply the reasons informing why people *do* act in a certain way; but in the meantime, they should supply the reasons informing why people *ought to* act in this way. Ideally, these reasons will coincide. When they do not, the game theorist should be careful to specify which type of explanation he is offering.

2. The Definition and Normative Role of Rationality

The concept of rationality is a hard nut to crack. Philosophers have spent ages failing to define it in a way that most deem satisfactory. Yet game theorists use the notion all of the time. By itself, that is not a problem, so long as they recognize what it means to do so.

In logic, we can discriminate between three types of definitions: lexical definitions, stipulative definitions, and analyses. Lexical definitions are those that are meant to describe the use of an existing, commonly used word. 'A *bachelor* is an *unmarried man*' qualifies as a lexical definition. A stipulative definition is one that proposes a meaning for a brand new word. 'A function is *differentiable* if and only if ...' is an example of a stipulative definition. And an analysis is a mixture of both. On the purpose of an analysis, Belnap [3] writes:

> Observe that in these cases the philosopher (theorist) neither intends simply to be reporting the existing usage of the community, nor would his or her purposes be satisfied by substituting some brand new word. In some cases it would seem that the philosopher's effort to explain the meaning of a word amounts to a proposal for a "good thing to mean by" the word.

Defining 'knowledge' as 'justified true belief' amounts to a good thing to mean by the word 'knowledge'.

By saying that a person is *rational*, the game theorist means that the person exhibits *optimizing behavior* with respect to her preferences, and that the person has *complete* *information*. Prima facia, this definition seems to be purely stipulative. However, it is often given the status of an analysis. That is highly problematic. When somebody uses a word that it is already in use, but they use it in some new way, it gives that person a way of sliding claims about that word in through the back door. Rationality has always been a word which carries normative force. To be more rational, in many instances, has simply meant to be more human[1]. It is what we might call a *strict good*. We would always prefer to have more than less. Thus, when a game theorist talks about rationality, he sometimes tricks people into thinking that he is able to tell them what they *ought to do*. But game theorists have no such power.

Because the notion of game-theoretic rationality is taken for granted, and because there is no external criterion by which we can measure which players in a game are more or less rational, the game-theoretic definition of rationality must not be treated as an analysis. In fact, along with this, there may be even further reason for not treating it as such. The assumption of game-theoretic rationality was originally drawn from Decision Theory, which analyzes the strategic behavior of single people in much the same way that game theory measures the behavior of multiple people at once. However, there is an often unacknowledged difference between those two situations. In single-player decision problems, decisions are completely removed from social considerations. In game theory problems, however, every player is made a member of a community. It is this community identity, I wish to propose, that poses a further problem for the game-theoretic definition of rationality.

Imagine the trivial decision problem in which I can decide between receiving a thousand dollars and not receiving a thousand dollars. I can even take just a portion of it, if I would like. Ceteris paribus, I have no reason for ever turning down the thousand dollars. I should take the whole thing. Now imagine the game theory problem in which someone is added to the game, while I am given the same choice. Now, the portion of the thousand dollars that I choose not to take will be given to the other person. According to most standard game theory models, the two player game is no less trivial than the other one. I should take every last dollar. But is this what I would actually do? No. Is it what I *ought* to do? No. This is because in the two player game, my decision has moved beyond one of pure ceteris paribus abstraction. Now, my actions imply explicitly-known consequences for somebody else. Insofar as I am now part of a community, I *ought* to take my neighbor into account. I should probably give him half. There is a lot of philosophical literature justifying the necessary role of normativity in rationality. I have just offered one argument.

Still, there are plenty of cases where economists conflate the criterion of efficacy with the criterion of prescriptiveness. Studies have been performed which indicate that people who study economics are significantly more likely to defect in social dilemmas than people who do not (see Frank, Gilovich, and Regan [6]). Why is this? There are two possible explanations: one draws on efficacy; the other, on prescriptiveness. Sorting out their differences becomes a question of the chicken or the egg. According to the first explanation, economists probably just defect more because they were predisposed to in the first place. Perhaps they got into the study of economics because they found the rational self-interest model to be appealing. According to the second, the studying of rational self-interest models has a way of convincing people that defecting is the right decision to make. Luckily, they also performed a study which showed that people receiving economics training

[1] In Aristotle's *Metaphysics*, he defines a 'man' as a 'rational animal'

that placed higher stress on rational self-interest were more likely to defect than people receiving economics training that was much less hawkish. This conclusion seems only to verify the latter explanation, and not the first.

Insofar as the game-theoretic concept of rationality is purely stipulative, it loses all of its normative force. At this point, the game theorist can no longer treat his models as giving categorical prescriptions. This does not prevent him, however, from giving hypothetical ones. Game Theory has only ever been a study of how players should optimize their behavior *given* that they have certain exogenously-determined goals. Essentially, this relegates the role of prescriptiveness to a matter of 'should'-statements invoking efficacy and hypothetical means-ends analysis[2].

3 The Inductive Approach

1. Inductive Evaluation

Most well-regarded solutions for the Prisoner's Dilemma indicate that the equilibrium strategy for players is to defect. Often, even, a solution will be respected only as long as it prescribes that strategy. In experimental cases, however, a sizable portion of subjects choose to cooperate. This is true for both (finitely) repeated and one-shot games (see Andreoni and Miller [1] and Clark and Sefton [4] on finitely repeated games, and Cooper, DeJong, Forsythe and Ross [5] on finitely repeated and one-shot games). While many game theorists are willing to accept this discrepancy, writing it off as this or that, I see it as a problem that deserves to be addressed. Furthermore, I think that it elucidates a more general problem within the practice of game theory itself. Game theorists sometimes are too willing to ignore inductive inferences in favor of deductive ones, even when the deductive ones yield unrealistic results.

Bertrand Russell [10] has the following to say about the study of mathematical formalizations:

> We tend to believe the premises because we can see that their consequences are true, instead of believing the consequences because we know the premises to be true. But the inferring of premises from consequences is the essence of induction; thus, the method of investigating the principles of mathematics is really an inductive method, and is substantially the same as the method of discovering general laws in any other science.

Empirical data can be viewed as providing some of the obvious consequences of game theory. That those consequences sometimes do not line up with an equilibrium suggests a drain on adequacy caused by the rationality assumption. No longer can game theorists ignore that disrepancy on prescriptive grounds, either. A purely stipulative definition of rationality forfeits that option. Rather than writing cooperators off as irrational, game theorists are now forced to address alternative notions of what constitutes rationality. The purpose will be to find a *good thing to mean* by it. It will also be to allow for more robust models.

[2]David Hume defines means-ends analysis as the practice of determining which course of action (the means) should be taken if a player wishes to satisfy a certain goal (the ends).

2. Inductive Generalization

Aside from offering us the opportunity to evaluate the use of our axioms, inductive styles of reasoning contribute to the establishment of axioms in the first place. Is it not always so that we first look out into world, perhaps solving a few examples, before we try to draw our assumptions? In practice, a game theorist already has his solutions ready. He is just searching for the set of axioms that will get him there. In reading Kreps and Wilson [9], this becomes abundantly clear. They state their objective from the beginning (p. 871):

> A strategy π should be such that for any information set h that is a singleton, player *i(h)* should not be able to change his strategy unilaterally and thereby improve this expected utility starting from h.

The resulting theory of sequential equilibria is contrived to yield those consequences. It is not the result of guessing at new possibilities for axioms and deductively discovering their conclusions. That process is too inefficient. Deduction is an end step, with induction at the beginning. New solution concepts are the result of searching for empirical counterexamples to earlier solution concepts, and trying to account for them in a better way.

Because the study of game theory naturally involves an inductive mode of arriving at generalizations, its practitioners should not be selective about what data they take seriously. Paying attention to the proportion of people who are willing to cooperate in Prisoner's Dilemma games is not only important, but necessary, for establishing the most adequate and efficacious theories that we can.

3. The Answer: Psychological Games

The answer to the adequacy problem, I suspect, lies in certain applications of psychological game theory. Psychological game theory is the branch of game theory that tries to account for the very sort of problem that I have so far been decrying. Originally devised by Geanakoplos, Pearce, and Stacchetti [7], psychological games are meant to account for the fact that human decision-making often not only depends upon the strategy of every player, but also on his or her beliefs. This allows utilities to depend on much more than just the material payoffs. It should be able to account for why certain people cooperate in the Prisoner's Dilemma. I realize that I am not the first to think this; nor am I even one of the first. However, I do believe that I have been able to justify more than most people why it is reasonable.

4 The Prisoner's Dilemma: Solution Concepts

1. Game Description

The Prisoner's Dilemma is a simultaneous game consisting of two players, each with two options: to Cooperate, or to Defect. While mutual cooperation is the Pareto efficient outcome (i.e., maximizes the sum of player's payoffs), the strategy Cooperate is strictly dominated by the strategy Defect. This induces the following normal form game

	c	d
C	α, α	β, δ
D	δ, β	γ, γ

with payoff structure $\delta > \alpha > \gamma > \beta, 2\alpha > \beta + \delta$.[3] These payoffs need not be merely material. However, in experiments, they almost inevitable are.

2. Equilibria

The Nash equilibrium for the Prisoner's Dilemma is the outcome at which neither player wishes to change his pure strategy. Since defecting strictly dominates cooperating, the only Nash equilibrium for the one-shot Prisoner's Dilemma is (D, d). Moreover, since there is only one strategy that remains after iteratively deleting strictly dominated strategies, the only mixed strategy equilibrium is the trivial $(\sigma(D) = 1, \sigma(d) = 1)$. This can be easily verified. Suppose that player 2 (column player) is using a mixed strategy, playing c with probability θ. At equilibrium, player 1 must be indifferent between choosing C or D. This yields

$$E_{u_1|\sigma_2}(C) = \theta\alpha + (1 - \theta)\gamma = \theta\beta + (1 - \theta)\delta = E_{u_1|\sigma_2}(D)$$

$$\rightarrow \theta = \frac{\delta - \gamma}{\alpha - \beta + \delta - \gamma}.$$

θ is only greater than zero when $\alpha + \delta > \beta + \gamma$. Even so, the only way achieve $\theta \leq 1$ is to have $\beta \leq \alpha$. That contradicts the payoff structure of the game itself. Therefore, there is no other mixed equilibrium in the Prisoner's Dilemma than to play Defect with probability 1. Because there are no proper subgames of the one-shot prisoner's dilemma, the subgame perfect equilibrium is trivial as well. It is the same as the Nash equilibrium. Therefore, none of the three traditional accounts of equilibrium are able to explain cooperation in the one-shot Prisoner's Dilemma. As it turns out, they are not able to explain it in the finitely-repeated version, either. All three predict that players will defect at every round.

Kreps and Wilson [9] offer a way of viewing equilibria that takes beliefs into account. It is called a sequential equilibrium (SE). Basically, it states that at a given information set, the player assigned to that set is playing optimally based upon his beliefs about the strategy of the other player. As long as those beliefs correspond to reality, there is a sequential equilibrium. Unfortunately, in the one-shot prisoner's dilemma, the sequential equilibrium notion, too, holds that both players should defect. That is because regardless of the probability that the other player is playing Cooperate, I can always maximize my expected payoff by defecting. The problem is so simple that the motivations enforcing the sequential equilibrium are also the motivations enforcing the mixed equilibrium, only now, in the sequential equilibrium, we can say that both players believe that the other will defect with probability one. In finitely repeated games, the sequential equilibrium

[3] Kreps *et al.* [8] point out that without the second condition, the strategy of both players cooperating at each stage is Pareto dominated by alternating between (C,d) and (D,c).

approach can sometimes explain why people cooperate (see Kreps, Milgrom, Roberts, and Wilson [8] on how this might occur). But why do they cooperate in the one-shot game?

Psychological game theorists have taken the notion of an SE one step further (see Geanakoplos, Pearce, and Stacchetti [7] and Battigalli and Dufwenberg [2] for two opposing models of doing this). They compute sequential equilibria, while allowing for modifications in the payoff structure that are due to certain psychological states. I think that this is the key for explaining why people cooperate in the one-shot Prisoner's Dilemma. In the next section, I will attempt to explain why some people would choose to cooperate. Because my purpose is only to demonstrate how inductive reasoning might be executed, my attempt will not be very robust. A more robust attempt is made by Cooper *et al.* [5] in "Cooperation without Reputation...".

3. Inductively Characterizing Types

For lack of data, the following must will be strictly hypothetical. However, it will draw on data from Cooper *et al* [5], in which players were shown to cooperate over 20% of the time in the one-shot Prisoner's Dilemma.

Suppose that an experiment has been performed in which subjects play the one-shot prisoners dilemma described in section 4.1. Here is the specific game:

	c	d
C	$9, 9$	$0, 10$
D	$10, 0$	$2, 2$

Suppose, further that, in the experiment, 25% of people cooperated while 75% defected. With just this most basic amount of information, one can inductively try to explain all of the players' actions. At the end of the game, players were asked whether they wish they would have changed their decisions based upon what they now know the other player chose. The model that I will be working with supposes a basic altruism that some players feel toward others within the community of the game. Suppose that each player has an altruism coefficient ϵ_i, so that the game now has the normal form

	c	d
C	$9+\epsilon_1, 9+\epsilon_2$	$\epsilon_1, 10$
D	$10, \epsilon_2$	$2, 2$

Suppose that each player also has a belief $\mu_i(\epsilon_{-i}) \sim N(\epsilon_i, \sigma^2)$ about how altruistic his opponent is. As you can see, that belief is based upon his own level of altruism. This means that if a person is average with regard to both his beliefs and his level of altruism, his beliefs about his opponent's level of altruism will equal his own level. This will be quite useful for setting up an equilibrium.

This model immediately separates players into three categories. *Pure altruists* are players who cooperate all of the time. They all exhibit $\epsilon > 2$. *Egoists* are players who never cooperate. They exhibit $\epsilon < 1$. And *mild altruists* are players whose level of altruism will sometimes let them be indifferent between choosing one pure strategy or the other. They exhibit $\epsilon \in [1, 2]$. It is clear that at equilibrium, pure altruists will always cooperate and egoists will always defect. However, at this point it is not so clear how mild altruists will act. Suppose that you and an opponent are both mild altruists with an equal level of altruism ϵ. At equilibrium, you will both believe that the other has $\epsilon_{-i} = \epsilon$. Otherwise, one of you will be able to profitably deviate from your equilibrium strategy. Thus, you both exhibit consistent beliefs (see Kreps and Wilson [9]). Meanwhile, if your opponent is playing Cooperate with probability θ, then you will be indifferent between choosing C or D. Therefore,

$$E[u_i(C)] = \theta(9+\epsilon) + (1-\theta)\epsilon = \theta(10) + (1-\theta)(2) = E[u_i(D)].$$

This yields $\theta = 2 - \epsilon$. Both players will play a mixed strategy in which they cooperate with probability $2-\epsilon$. It is important to note that this represents an equilibrium between two perfectly average mild altruists. If a mild altruist encounters an egoist at equilibrium, then he will want to defect (since $\epsilon < 2$), whereas if he encounters a pure altruist, then he will want to cooperate (since $\epsilon > 1$).

Now, suppose that when players were asked whether they would like to have changed their decision, nobody who ended up in a symmetric situation said yes. However, 80% of those who cooperated while their opponent defected said yes, and 20% of those who defected while their opponent cooperated also said yes. Essentially, this is a way of asking players whether they had reached an equilibrium. It might sound odd to say that someone who defected might want to have cooperated after seeing that his opponent did, but that is not so. It is easy to conceive of a player who, expecting his opponent to choose defect, also chose it, and then regretted his decision once he realized that the other player was cooperating. This would also be consistent with the way that mild altruists behave with pure altruists at equilibrium. Given the supposed data, I would conclude that everyone who made such a request had to be a mild altruist. Egoists and pure altruists would never want to change their decisions, because they both have dominant strategies. Since presumably a player's choice is independent of his opponent's, we can conclude that the proportion of mild altruists will be uniform across opponent choices. Therefore, we can conclude that 80% of the players who ended up in (C,c) were mild altruists, while 20% of the players who ended up in (D,d) were. Therefore, we can conclude that 80% of all cooperators and 20% of defectors are mild altruists. Applying this result, and using the fact that 25% of people cooperated overall, we can determine what proportion of players falls under each category.

$$p(\text{mild altruist}) = \frac{4}{5} \times \frac{1}{4} + \frac{1}{5} \times \frac{3}{4} = \frac{7}{20},$$

$$p(\text{pure altruist}) = \frac{1}{5} \times \frac{1}{4} = \frac{1}{20},$$

$$\text{and } p(\text{egoist}) = \frac{4}{5} \times \frac{3}{4} = \frac{12}{20}.$$

Since we know that 25% of people cooperated, and we also know that pure altruists always cooperated while egoists never did, we can now determine the average probability with which a mild altruist cooperated.

$$p(C) = \frac{12}{20} \times 0 + \frac{1}{20} \times 1 + \frac{7}{20} \times \theta = .25,$$

which yields $\theta = \frac{4}{7}$. This is also implies that at the equilibrium mentioned above, in which both players are perfectly average among mild altruists, we can conclude that the average altruism level for mild altruists was $2 - \frac{4}{7} = \frac{10}{7}$. This is a very helpful result. Presumably, if we were able to conduct many experiments, where we would alter the payoffs only slightly from one experiment to the next, and where we would compare the results of each experiment to the payoffs within, it might soon be possible to generate a regression function that estimates people's levels of altruism. Upon inputting the payoffs themselves, we would then be able to predict with a high level of accuracy how players will behave in the Prisoner's Dilemma. We would also then be able to test the adequacy of the altruism theory that I have just proposed.

Notice the strengths of the inductive method. Using a model that treats people as receiving some extra benefit from cooperating, I have been able to provide a possible explanation for why they would. I have even been able to pin down some important features of the altruism coefficient. I can now provide prescriptions to players who exhibit a certain level of altruism, capturing the necessarily hypothetical nature of game-theoretic prescriptiveness.

References

[1] J. Andreoni and J.H. Miller. "Rational Cooperation in the Finitely Repeated Prisoner's Dilemma: Experimental Evidence", *The Economic Journal* 103 (1993).

[2] P. Battigalli and M. Dufwenberg. "Dynamic Psychological Games", *Journal of Economic Theory* 144 (2009).

[3] N. Belnap. *Notes on the Art of Logic*, University of Pittsburgh, 2009. *unpublished*

[4] K. Clark and M. Sefton. "The Sequential Prisoner's Dilemma: Evidence on Reciprocation", *The Economic Journal* 111 (2001).

[5] R. Cooper, D.V. DeJong, and R. Forsythe. "Cooperation without Reputation: Experimental Evidence from Prisoner's Dilemma Games", *Games and Economic Behavior* 12 (1996).

[6] R.H. Frank, T. Gilovich, and D.T. Regan. "Does Studying Economics Inhibit Cooperation?", *The Journal of Economic Perspectives* 7, (1993).

[7] J. Geanakoplos, D. Pearce, and E. Stacchetti. "Psychological Games and Sequential Rationality", *Games and Economic Behavior* 1 (1989).

[8] D. Kreps, P. Milgrom, J. Roberts, R. Wilson. "Rational Cooperation in the Finitely Repeated Prisoner's Dilemma", *Journal of Economic Theory* 27 (1982).

[9] D. Kreps and R. Wilson. "Sequential Equilibria", *Econometrica* 50 (1982).

[10] B. Russell. "The Regressive Method of Discovering the Premises of Mathematics", *Mathematics, Matter and Method: Philosophical Papers, Vol. 1*, Cambridge University Press, Cambridge, 1907.

[11] K. Ritzberger. *Foundations of Non-Cooperative Game Theory*, Oxford University Press, New York, 2002.