



LUDWIG-
MAXIMILIANS-
UNIVERSITÄT
MÜNCHEN

VOLKSWIRTSCHAFTLICHE FAKULTÄT



Gallice, Andrea:

Best Responding to What? A Behavioral Approach to One Shot Play in 2x2 Games

Munich Discussion Paper No. 2007-6

Department of Economics
University of Munich

Volkswirtschaftliche Fakultät
Ludwig-Maximilians-Universität München

Online at <https://doi.org/10.5282/ubm/epub.1365>

Best Responding to What?

A Behavioral Approach to One Shot Play in 2x2 Games

Andrea Gallice*

University of Munich
and
European University Institute

Abstract

We introduce a simple procedure to be used for selecting the strategies most likely to be played by inexperienced agents who interact in one shot 2x2 games. We start with an axiomatic description of a function that may capture players' beliefs. Various proposals connected with the concept of mixed strategy Nash equilibrium do not match this description. On the other hand minimax regret obeys all the axioms. Therefore we use minimax regret to approximate players' beliefs and we let players best respond to these conjectured beliefs. When compared with existing experimental evidences about one shot matching pennies games, this procedure correctly indicates the choices of the vast majority of the players. Applications to other classes of games are also explored.

Keywords: prediction, beliefs, mixed strategy Nash equilibrium, minimax regret, matching pennies, experiments.

JEL classification: C72, C91.

1 Introduction

How should an inexperienced agent who is facing a simple strategic situation, like for instance a one shot 2x2 game, decide which strategy to play? A rational, or Bayesian, approach for dealing with such a decision under uncertainty (in a spirit similar to Savage, 1954) would be the following:

*Contact: Seminar for Economic Theory, Ludwig Maximilians University (LMU), Ludwigstrasse 28 (Rg), 80539 Munich. Email: andrea.gallice@eui.eu.

- 1 - the player forms a belief about what his opponent will play.
- 2 - the player chooses the strategy which best responds to this belief.

Interpreting it as a heuristics and acknowledging players' heterogeneity, this rule cannot be expected to describe the behavior of every single agent. Still some recent papers confirm that in simple strategic interactions the majority of individuals behave coherently with their beliefs.¹ Nyarko and Schotter (2002) study a 60 times repeated 2x2 game and find that around 75% of the players do indeed best respond to their stated beliefs. For the case of 3x3 games, Rey-Biel (2004) considers 10 one shot games and finds a similar rate of compliance while 55% is the percentage found by Costa-Gomes and Weizsäcker (2005) using data about 14 (more complex) one shot games. It is therefore a conservative guess to expect that, in one shot 2x2 games, at least half of the individuals play consistently with their beliefs. We want to capture the behavior of this majority of players.

In order to do so we initially focus on the process of beliefs formation. The goal is to find a function that may approximate players' beliefs without the need of having to explicitly elicit them. Using an axiomatic approach we first list a limited number of desirable properties that, in accordance with experimental studies and behavioral regularities, a belief function must fulfill. Then, we check some existing concepts commonly used in game theory and decision theory to see which of them, if any, fulfills all the requirements.

We first show that, contrary to a very authoritative strand of research (starting with Aumann, 1987, who simplified Harsanyi, 1973), the concept of mixed strategy Nash equilibrium is not a suitable candidate for approximating players' beliefs. Indeed mixed equilibria are still problematic both in theory and in practice (see the literature review). In particular some features of mixed equilibria clearly contradict our axiomatic description.

At the opposite we find minimax regret (a criterion for dealing with choices under uncertainty originally proposed by Savage, 1951) to be the unique candidate to obey all the axioms. Therefore we propose minimax regret as a proxy for players' beliefs and we claim that the majority of players play "as if" they were best responding to these approximated beliefs.

This conjecture is tested in the second part of the paper. The predictions stemming from the procedure (best respond to beliefs equal to the minimax regret distribution of the opponent) are compared with experimental evidences about different versions of 2x2 one shot matching pennies games. To forecast players' choices in this class of games can be particularly problematic because the Nash indication is often misleading (see for instance Ochs, 1995 and Goeree and Holt, 2001). Our procedure proves to be an effective way to identify the strategies which are more likely to be

¹These papers elicit players' beliefs using a proper quadratic scoring rule, such that for the players "telling the truth" is optimal.

played. In fact it correctly predicts the actual choices of around 80% of the players.

We also apply the procedure to other classes of 2x2 games and we analyze its relationship with the Nash prediction. An interesting result is that the procedure selects a single outcome even in games that have multiple Nash equilibria such that it contributes to the debate on equilibrium selection (see Straub, 1995 and Haruvy and Stahl, 2004).

A number of studies that focus on how people play one shot simultaneous games and investigate the issue of beliefs formation are closely related to this paper. For instance Stahl and Wilson (1995), Costa-Gomes et al. (2000) and Camerer et al. (2004) hypothesize and test the existence and relative importance of various archetypes of players that differ in the prior they have about the degree of sophistication of their opponents. An important difference with respect to these papers is that we do not consider heterogeneity. We recognize heterogeneity to be a very important feature of human behavior but still here we favor simplicity. Therefore we only model the behavior of the majority of players. Our contribution is to provide a fast, easy to implement and effective procedure that gives rule of thumb predictions about the outcome one should expect to arise in simple games.

The paper is structured as follows: Section 2 briefly reviews the theoretical interpretations and the empirical relevance of mixed strategy Nash equilibria. Section 3 presents the basic axioms that must be satisfied by a belief function. Section 4 checks the compliance to these axioms of various candidate functions and it shows that minimax regret is the unique one to obey them all. Section 5 formally defines the procedure to be used to predict the outcome of simple games. These predictions are then tested in Section 6. Section 7 concludes.

2 Literature review on mixed equilibria

The idea that agents may randomize over a set of actions (i.e. they may use mixed strategies) dates back to Borel (1921) and was then enriched and developed by Von Neumann and Morgenstern (1944). The concept of Nash equilibrium appeared a few years after (Nash, 1951) and rapidly became the most important solution concept of game theory. In his paper Nash also presented the concept of mixed strategy Nash equilibrium, a notion that led to the famous theorem about the existence of a Nash equilibrium in any finite game.

In a mixed strategy Nash equilibrium (shortened in what follows to *MSNE*) each player randomizes over (some of) his pure strategies according to a probability distribution that makes the other players indifferent about what to play. Therefore no player has any strict incentive to deviate and the mixed strategy profile identifies a Nash equilibrium of the game.

The problem to look for this specific probability distribution shares quite many features with

the following problem of elementary physics.

Exercise: You have a rigid rod of negligible weight which is 1 meter long. On one side of it there is a weight of 1 Kg, on the other side a weight of 4 kg. Where does the rod have to be pivoted in order for it to be in equilibrium?




Figure 1: the equilibrium problem.

The rod is in equilibrium if $4l_1 = l_2$ where $l_1 + l_2 = 1$ and l_1 is the distance from the pivot to the weight of $4Kg$ and l_2 is the distance from the pivot to the weight of $1Kg$. Therefore $l_1 = 0.2m$ and $l_2 = 0.8m$. The pivot has to be put closer to the heavier side. Indeed $\frac{4}{5}$ of the rod have to be given to the lighter side in order to impose the equilibrium.

A similar situation arises when we compute *MSNE*: the equilibrium condition among pure strategies is imposed and then, as a consequence, the appropriate probability distribution is retrieved. And as it happens in the rod example, mixed equilibria often allocate larger probabilities to “lighter” strategies, i.e. to strategies that are associated with lower payoffs.² This feature undermines the predictive power of *MSNE*. Another unappealing feature of mixed equilibria is the so called “no own payoff effect” which will be shortly discussed.

To sum up *MSNE* often appear as quite an artificial construction and indeed they are still problematic both in theory and in practice. To quote Rubinstein (1991): “*The concept of mixed strategy has often come under heavy fire*”. On the theoretical side different interpretations have been given about mixed equilibria and a general consensus is still missing.³ On the empirical side the relevance of *MSNE* in capturing agents’ behavior has also been heavily questioned.

²The following game exactly mimics the rod example. And indeed the mixed equilibrium is given by $(\frac{1}{5}T + \frac{4}{5}B)$ for both players.

	<i>T</i>	<i>B</i>
<i>T</i>	4, 4	0, 0
<i>B</i>	0, 0	1, 1

³See for instance Osborne and Rubinstein (1994), Section 3.2.

2.1 Theoretical interpretations of mixed equilibria

Taken at face value a mixed strategy Nash equilibrium prescribes a player to select the strategy to play according to a specific probability distribution. In other words players should deliberately randomize. This view may make sense when players are interacting in repeated games and they do not want their pattern of play to be predictable.⁴ But in one shot interactions such a randomizing behavior looks less realistic.

A second interpretation considers the *MSNE* probabilities as indicating the steady state frequencies of pure strategies when games are played in large populations (see for instance Rosenthal, 1979). Each player chooses a pure strategy but in the entire population the *MSNE* distribution should emerge. Again this interpretation does not look appropriate in the context we are interested in, namely a one shot interaction between two players.

Very much related with the ideas presented in this paper are the interpretations that link *MSNE* with the beliefs of the players. Harsanyi (1973) presented the so called purification interpretation of *MSNE*. The idea is that players play pure strategies. The twist is that each player's choice is based on some private information. This means that player i knows what to play but the information sets of i 's opponents are not enough precise to allow them to be sure about i 's choice. *MSNE* captures this uncertainty (or ignorance).

Aumann (1987) further simplified this idea. The claim (see also Aumann and Brandenburger, 1995 and Reny and Robson, 2003) is that, even in the absence of this small amount of private information, players are still unsure about the opponents' moves. Therefore the probability distribution that the *MSNE* assigns to player i can be directly interpreted as the beliefs all the other players hold about i 's choice. We will come back to this point in Section 4.1.

2.2 Empirical relevance of mixed equilibria

Because of the existence of all these different theoretical interpretations, it is no surprise that many experimental studies have been designed with the aim of testing the empirical relevance of *MSNE*. In particular researchers focused on the study of matching pennies (MP) games, i.e. games that have a unique Nash equilibrium which is in mixed strategies.⁵ The common design of these experiments consisted in letting subjects repeatedly play the same version of a MP game. Results show (for a detailed review see also Camerer, 2003) that long run frequencies of pure strategies are not too far from the *MSNE* predictions. Still players' behavior at an individual

⁴Evidences of such a mixing behavior appear in professional sports. See for instance Walker and Wooders (2001), Chiappori et al. (2002) and Palacios-Huerta (2003).

⁵Important contributions are Mookherjee and Sopher (1994), Ochs (1995), McKelvey and Palfrey (1995), Erev and Roth (1998), McKelvey et al. (2000), Tang (2001) and Goeree et al. (2003).

level is often significantly different from the one indicated by the equilibrium.

A different, and less investigated, question is to study how agents behave in front of a single interaction.⁶ In a one shot MP game players cannot learn over time and the incentive to maintain an unpredictable pattern of play simply does not exist. In these cases the *MSNE* prediction is useless. In fact the equilibrium indicates a mixed distribution while the result of the game will be a specific outcome given by the intersection of two pure strategies.

Our paper is focused on one shot games mainly for three reasons: first, we claim our theory to be able to capture the behavior of inexperienced players; second, as just mentioned, one shot individuals' play has been less investigated; third, we think that enough real life situations are more likely to be similar to one off events rather than to repeated interactions.

3 An axiomatic approach to model players' beliefs in 2x2 games

We present some axioms that, in our view, capture some very basic properties of how the beliefs of inexperienced and boundedly rational players should look like. These axioms are quite general and therefore, perhaps unfortunately, they do not characterize a single belief function. Still they are enough to clearly discriminate among various functions that are commonly used in game theory and in decision theory. In fact, in order to find a function that may approximate players' beliefs, we do not want to introduce new ad hoc formulas. At the opposite we rely on concepts that already exist and that are widely used, though possibly for different purposes. Given that the conjectured belief function will be used for predictive purposes, we perceive this to be a useful result.

Consider the following game, where player $i \in \{A, B\}$ can choose between strategies H_i and T_i .

		H_B	T_B
1)	β_{H_A}	H_A	$a, \cdot \quad b, \cdot$
	β_{T_A}	T_A	$c, \cdot \quad d, \cdot$

We model the beliefs player B holds about what player A will play. This is the reason why the payoff matrix is incomplete and only the payoffs of player A appear. Whilst keeping in mind the example of an inexperienced boundedly rational player, B 's beliefs are considered as being just a function of player A 's payoffs. This may look like a drastic requirement but it can be seen as a simplification of the fact that B 's beliefs are mainly a function of A 's payoffs and then, as

⁶A notable exception is Goeree and Holt (2004) that presents a model of iterated noisy introspection for one shot interactions which is then tested over a large number of games.

a secondary effect that we do not consider, B 's beliefs may also be influenced by higher order speculations that involve B 's payoffs.

The focus on B 's beliefs implies no loss of generality given that a similar analysis can be done for what concerns A ' beliefs. We indicate with:

- β_{H_A} the beliefs player B hold about player A playing strategy H_A
- β_{T_A} the beliefs player B hold about player A playing strategy T_A

Finally we do not put any particular restriction on the values of the payoffs a , b , c and d . According to us, player B 's belief function must obey to the following five axioms:

[A1] Functional form

$$\beta_{H_A} = \beta_{H_A}(a, b, c, d), \beta_{T_A} = \beta_{T_A}(a, b, c, d)$$

[A2] Consistency with probability distribution

$$\beta_{H_A} \geq 0, \beta_{T_A} \geq 0 \text{ and } \beta_{H_A} + \beta_{T_A} = 1 \text{ for any } a, b, c, d.$$

The first axiom simply formalizes what we already mentioned, i.e. the fact that B 's beliefs have to be a function of the payoffs of player A . This implies that player B realizes A does respond to changes in his own payoffs. This own payoff effect is a very robust feature of games played in experiments (for clear evidences of this effect see, among others, Ochs, 1995 and Goeree and Holt, 2001). The second axiom requires the function to identify a proper probability distribution, i.e. beliefs that are non negative and that sum up to one. Notice that, because of the relation $\beta_{T_A} + \beta_{H_A} = 1$, a single belief is enough to define the entire distribution. Therefore, from now on, we focus on β_{H_A} .

[A3] Continuity

β_{H_A} is continuous in all its arguments.

[A4] Monotonicity

β_{H_A} is weakly increasing in a and b . β_{H_A} is weakly decreasing in c and d .

[A5] Consistency with rows or columns switch

Continuity (A3) is required since there are no evident reasons for B 's beliefs to jump in a discrete way given small changes in the arguments of the function. The monotonicity axiom (A4)

defines the sign of the already mentioned own payoff effect and it is in line with a large amount of experimental evidences (among others Ochs, 1995; Goeree and Holt, 2001; Goeree et al., 2003). The axiom states that players believe their opponent to be more likely to play strategies that look better. In other words if the payoffs associated with strategy H_A increase then the probability that player B assigns to the event of A playing that strategy cannot decrease.

The last axiom indicates that the beliefs of player B have to consistently react to the payoff structure of the game and they do not depend on the labelling of the strategies. To better understand the meaning of this axiom consider the two games that are reported below. With respect to the original game (Game 1 above), Game 1' is such that the payoffs of the two rows have been switched while Game 1'' is such that the payoffs of the two columns have been switched. Axiom 5 requires that, whenever $a \neq c$ and $b \neq d$, the probability that B 's belief function allocates to the event of A playing strategy H_A in Game 1 must be equal to the probability that the same function allocates to the event of A playing T_A in Game 1'. In other words $\beta_{H_A} = \beta'_{T_A}$. At the opposite the belief function must allocate the same probabilities to the events of A playing H_A in Game 1 and in Game 1'' such that $\beta_{H_A} = \beta''_{H_A}$. In fact a column switch does not change the relative attractiveness of strategies H_A and T_A from player A 's point of view.

	<table style="border-collapse: collapse; margin-left: auto; margin-right: auto;"> <tr> <td style="padding: 0 10px;"></td> <td style="padding: 0 10px;"></td> <td style="padding: 0 10px; text-align: center;">H_B</td> <td style="padding: 0 10px; text-align: center;">T_B</td> </tr> <tr> <td style="padding: 0 10px; vertical-align: middle;">1')</td> <td style="padding: 0 10px; vertical-align: middle;">β'_{H_A}</td> <td style="padding: 0 10px; vertical-align: middle;">H_A</td> <td style="padding: 0 10px; vertical-align: middle;">c, \cdot</td> </tr> <tr> <td style="padding: 0 10px;"></td> <td style="padding: 0 10px; vertical-align: middle;">β'_{T_A}</td> <td style="padding: 0 10px; vertical-align: middle;">T_A</td> <td style="padding: 0 10px; vertical-align: middle;">a, \cdot</td> </tr> </table>			H_B	T_B	1')	β'_{H_A}	H_A	c, \cdot		β'_{T_A}	T_A	a, \cdot		<table style="border-collapse: collapse; margin-left: auto; margin-right: auto;"> <tr> <td style="padding: 0 10px;"></td> <td style="padding: 0 10px;"></td> <td style="padding: 0 10px; text-align: center;">H_B</td> <td style="padding: 0 10px; text-align: center;">T_B</td> </tr> <tr> <td style="padding: 0 10px; vertical-align: middle;">1'')</td> <td style="padding: 0 10px; vertical-align: middle;">β''_{H_A}</td> <td style="padding: 0 10px; vertical-align: middle;">H_A</td> <td style="padding: 0 10px; vertical-align: middle;">b, \cdot</td> </tr> <tr> <td style="padding: 0 10px;"></td> <td style="padding: 0 10px; vertical-align: middle;">β''_{T_A}</td> <td style="padding: 0 10px; vertical-align: middle;">T_A</td> <td style="padding: 0 10px; vertical-align: middle;">d, \cdot</td> </tr> </table>			H_B	T_B	1'')	β''_{H_A}	H_A	b, \cdot		β''_{T_A}	T_A	d, \cdot
		H_B	T_B																								
1')	β'_{H_A}	H_A	c, \cdot																								
	β'_{T_A}	T_A	a, \cdot																								
		H_B	T_B																								
1'')	β''_{H_A}	H_A	b, \cdot																								
	β''_{T_A}	T_A	d, \cdot																								

4 The candidate functions

In this section we present various candidate functions and we check their compliance to the axioms. We first analyze a number of proposals connected with the concept of mixed strategy Nash equilibrium. Then we turn our attention to the minimax regret, a criterion normally used in decision theory. In line with the axiomatic analysis of the previous section, we check how the following functions perform in approximating the beliefs player B holds about what player A will play in Game 1.

4.1 Mixed strategy Nash equilibrium of player A

As mentioned in the literature review, the interpretation that mixed strategy equilibria could capture players' beliefs dates back to Harsanyi (1973) and it was later simplified by Aumann (1987). According to Aumann's view (see also Aumann and Brandenburger, 1995 and Reny and

Robson, 2003) the probability distribution that the *MSNE* attaches to player i represents the common belief all the other players have about what i will play.

In the context of our 2x2 game this boils down to the hypothesis that A 's equilibrium distribution captures B 's beliefs. This interpretation suffers from one fundamental drawback. By construction the *MSNE* of player A is exclusively a function of the payoffs of player B and it is totally unrelated with A 's payoff. In the literature this unappealing feature of mixed equilibria is called “no own payoff effect”. To make things clear consider Game 2 where, with respect to Game 1, also the payoffs of player B appear.

		H_B	T_B
2)	β_{H_A}	H_A	a, x
	β_{T_A}	T_A	c, w
			b, y
			d, z

The distribution that the *MSNE* assigns to player A is defined by the probability \tilde{p} that makes player B indifferent between playing H_B or T_B . Therefore \tilde{p} solves $\tilde{p}x + (1 - \tilde{p})w = \tilde{p}y + (1 - \tilde{p})z$ such that $\tilde{p} = \frac{z-w}{x-y+z-w}$ and the equilibrium component of player A is given by $(\tilde{p}H_A + (1 - \tilde{p})T_A)$. Interpreting this probability distribution as B 's beliefs we would have:

$$\beta_{H_A} = \frac{z - w}{x - y + z - w}$$

Notice the counter intuitive implications of this proposal: B 's beliefs remain constant no matter how A 's payoffs may change. For instance $\beta_{H_A} = \tilde{p}$ in the game above as well as in a similar game where the payoff a (or b, c, d) is substituted with, let's say, $5a$. This proposal therefore fails Axiom 1 and Axiom 4. Beliefs of this kind would not capture any own payoff effect. This is a serious limitation given that the existence and the importance of such an effect is testified by many experimental studies.

Notice furthermore that the predictive power of this proposal would be null. In fact any strategy (or combinations of strategies) in the support of the *MSNE* of player B is a best response to these conjectured beliefs.

4.2 Mixed strategy Nash equilibrium of player B

The analysis of the previous section may suggest the use of the *MSNE* of player B as a way to approximate B 's beliefs. Indeed the probabilities implied by the *MSNE* of B are a function of A 's payoffs. More precisely the *MSNE* of B is given by $(\tilde{q}H_B + (1 - \tilde{q})T_B)$ where \tilde{q} solves

$\tilde{q}a + (1 - \tilde{q})b = \tilde{q}c + (1 - \tilde{q})d$ such that $\tilde{q} = \frac{d-b}{a-c+d-b}$ and therefore, under this proposal, we would have:

$$\beta_{H_A} = \frac{d-b}{a-c+d-b}$$

The problem with this formulation is that beliefs of this kind do not obey to the monotonicity requirement. In fact β_{H_A} is decreasing in a and b and increasing in c and d while Axiom 4 requires the opposite behavior. Assume that payoffs a and b increase such as to make strategy H_A more attractive from A 's point of view. The probability weight that the *MSNE* attaches to the correspondent strategy H_B will get smaller (the intuition is the same provided by the rod example presented in Section 2). Therefore the mixed equilibrium of player B does capture the own payoff effect but with the wrong sign.

Trying to improve on the limits of the last proposal one may be tempted to approximate B 's beliefs by switching the probabilities implied by the *MSNE* of B . In other words by setting $\beta_{H_A} = 1 - \tilde{q}$ such that:

$$\beta_{H_A} = \frac{a-c}{a-c+d-b}$$

This proposal is a function of A 's payoffs and it now satisfies the monotonicity axiom. Still it fails Axiom 5. In fact, following this conjecture, player B should keep the same beliefs ($\beta_{H_A} = \beta'_{H_A}$) also in game $1'$ (see before) where again the *MSNE* of B is given by $\tilde{q} = \frac{d-b}{a-c+d-b}$ such that $1 - \tilde{q} = \beta'_{H_A}$. This would violate A5 whenever $a \neq c$ and $b \neq d$.

4.3 A new proposal: the minimax regret

We now present a new and unusual candidate for approximating players' beliefs. This proposal is based on an instrumental use of the minimax regret criterion. More precisely we claim that the beliefs of player j about what player i will play can be approximated by the minimax regret of player i .

Minimax regret, originally proposed by Savage (1951), is a concept which found its main applications as a selection criterion in decision theory (starting with Milnor, 1954).⁷ The minimax regret criterion prescribes a player who has to make a decision under uncertainty to choose the

⁷More recently minimax regret has been used to model the behavior of subjects with limited rationality (Bergemann and Schlag, 2005) as well as a way to deal with missing data in econometrics (Manski, 2005) and it also appears in the artificial intelligence literature (Brafman and Tennenholtz, 2000).

action that minimizes the maximum regret he may suffer. The regret of player i is defined as the difference between the best payoff i could have got if he knew what his opponent (another player or Nature) had played and the payoff player i actually got.

The first step to properly compute the minimax regret consists in building the regret matrix which captures these differences. In the context of Game 1 (reproduced below) one needs to know which is the largest payoff between a and c and between b and d . Let us assume that $a > c$ and $d > b$.⁸ In this case the regret matrix is given by R_1 :

$$\begin{array}{c}
 \begin{array}{cc|cc}
 & & H_B & T_B \\
 \hline
 1) \beta_{H_A} & H_A & a, \cdot & b, \cdot \\
 \beta_{T_A} & T_A & c, \cdot & d, \cdot \\
 \hline
 \end{array}
 &
 R_1) &
 \begin{array}{cc|cc}
 & & H_B & T_B \\
 \hline
 H_A & 0, \cdot & d - b, \cdot & (p_r) \\
 T_A & a - c, \cdot & 0, \cdot & (1 - p_r) \\
 \hline
 \end{array}
 \end{array}$$

Strategy H_A attains minimax regret if $a - c > d - b$ while strategy T_A attains minimax regret with a reversed relation. Taking this specification as a belief function would clearly be unsatisfactory given that such a proposal would fail the continuity axiom. The use of mixed strategies solves this problem.

The mixed minimax regret is defined by the probability distribution (identified by \tilde{p}_r , where the index r indicates regret) that equalizes the expected regret of the two strategies. This optimal \tilde{p}_r solves $\tilde{p}_r(d - b) = (1 - \tilde{p}_r)(a - c)$ so that $\tilde{p}_r = \frac{a-c}{a-c+d-b}$. According to the conjecture of this paper $\beta_{H_A} = \tilde{p}_r$ should hold and thus:

$$\beta_{H_A} = \frac{a - c}{a - c + d - b}$$

Once again, referring to Game 1 above, this means that player B approximately believes player A will play strategy H_A with probability $\frac{a-c}{a-c+d-b}$ and strategy T_A with probability $\frac{d-b}{a-c+d-b}$.⁹

It is easy to show that this candidate function obeys all the axioms. In fact $\beta_{H_A} = \beta_{H_A}(a, b, c, d)$ (A1), $\beta_{H_A} \geq 0$, $\beta_{T_A} \geq 0$ and $\beta_{H_A} + \beta_{T_A} = 1$ (A2), β_{H_A} is continuous in its arguments (A3) and it also fulfills the monotonicity requirement (A4).¹⁰ Finally it also obeys Axiom 5. In fact, referring to games 1' and 1'' presented before, we have that $\tilde{p}_r = \frac{d-b}{a-c+d-b}$ in Game 1' such that $\beta_{H_A} = \beta'_{T_A}$ and $\tilde{p}_r = \frac{a-c}{a-c+d-b}$ in Game 1'' such that $\beta_{H_A} = \beta''_{H_A}$.

⁸This assumption implies no loss of generality. It is in fact easy to show that the minimax regret proposal satisfy the axioms also in the case of $a < c$ and $d < b$ as well as in the cases with weakly or strictly dominated strategies.

⁹These probabilities are analogous to the ones identified by the $\beta_{H_A} = 1 - \tilde{q}$ proposal. Still the latter fails Axiom 5 while a belief function based on minimax regret fulfills this requirement. More in general in any 2x2 game where a non degenerated *MSNE* exists, its probability distribution is either the same or the mirror image of the minimax regret distribution of the other player (see Gallice, 2006b).

¹⁰First partial derivatives of β_{H_A} with respect to its arguments are given by: $\frac{\partial \beta_{H_A}}{\partial a} = -\frac{\partial \beta_{H_A}}{\partial c} = \frac{d-b}{(a-b+d-c)^2} \geq 0$ and $\frac{\partial \beta_{H_A}}{\partial b} = -\frac{\partial \beta_{H_A}}{\partial d} = \frac{a-c}{(a-b+d-c)^2} \geq 0$. Second derivatives show that β_{H_A} is concave in a and b and convex in c and d .

4.4 A summary and an example

The following table summarizes how the various proposals that we just discussed perform in matching the five axioms¹¹. Axioms are identified as: functional form (A1), consistency with probability distribution (A2), continuity (A3), monotonicity (A4) and consistency with rows or columns switch (A5).

candidates \ axioms	A1	A2	A3	A4	A5
<i>MSNE</i> of pl. <i>A</i>	Y	N	Y	N	Y
<i>MSNE</i> of pl. <i>B</i>	Y	Y	Y	N	N
1 - <i>MSNE</i> of pl. <i>B</i>	Y	Y	Y	Y	N
Pure minimax regret of pl. <i>A</i>	Y	Y	N	Y	Y
Mixed minimax regret of pl. <i>A</i>	Y	Y	Y	Y	Y

Table 1: compliance to the axioms of the candidate functions.

Among the various candidates, the (mixed) minimax regret proposal is the unique one to satisfy all the axioms and thus the unique one to qualify for approximating players' beliefs. Before properly defining a procedure which is based on these conjectured minimax regret beliefs, we present a useful example. This example is meant to show how simple is the process of approximating beliefs using minimax regret, under lining once more the inadequacy of proposals connected with the *MSNE* concept.

Consider Game 3 where $k \in (-\infty, \infty)$ such that the game encompasses the cases of a matching pennies game (for $k \in (-1, \infty)$) and of a game with a dominant strategy (for $k \in (-\infty, -1]$).

$$\begin{array}{c}
 \begin{array}{cc} & \begin{array}{cc} \hline H_B & T_B \\ \hline \end{array} \\
 3) \quad \beta_{H_A} & \begin{array}{cc} H_A & k, -1 \\ & -1, 1 \end{array} \\
 \beta_{T_A} & \begin{array}{cc} T_A & -1, 1 \\ & 1, -1 \end{array} \\
 \hline
 \end{array} \\
 \\
 \begin{array}{cc}
 \begin{array}{cc} \hline H_B & T_B \\ \hline \end{array} & \begin{array}{cc} \hline H_B & T_B \\ \hline \end{array} \\
 R_1 = \begin{array}{cc} H_A & 0, 2 \\ T_A & k + 1, 0 \end{array} & R_2 = \begin{array}{cc} H_A & -1 - k, 2 \\ T_A & 0, 0 \end{array} \\
 k \in (-1, \infty) & k \in [-\infty, -1)
 \end{array}
 \end{array}$$

For $k \in (-1, \infty)$ the regret matrix is given by R_1 . The minimax regret of player *A* is given by $(\tilde{p}_r H_A + (1 - \tilde{p}_r) T_A)$ where $\tilde{p}_r = \frac{1+k}{3+k}$. Following our proposal $\beta_{H_A} = \frac{1+k}{3+k}$ captures *B*'s beliefs

¹¹Gallice (2006a) considers a richer set of axioms but the results of the analysis are analogous.

about A playing strategy H_A . This function appears as the bold concave curve in Figure 2 which focuses on the beliefs of player B about what A will play.

For $k \in (-\infty, -1]$ the regret matrix is given by R_2 such that $\tilde{p}_r = \beta_{H_A} = 0$, i.e. player B attaches a null probability to the event of A playing H_A . In Figure 2 this appears as the bold line that lies on the horizontal axis for $k \leq -1$.

The other two functions (dashed lines) depict, respectively, the probability that the mixed strategy Nash equilibrium assigns to player A playing strategy H_A ($\frac{1}{2}$) and to player B playing strategy H_B ($\frac{2}{3+k}$). The figure thus highlights the problems which were mentioned in the previous section: the *MSNE* of player A does not respond to a change in A 's payoff while the *MSNE* of player B does respond to a change in k but not in the desired direction. Notice that the functions for the minimax regret and for the *MSNE* of the two players intersect just once. The intersection happens for the unique k (in this case $k = 1$, symmetric game) for which all the three functions reach a value of $\frac{1}{2}$.




Figure 2: beliefs approximation through minimax regret in Game 3.

5 The procedure

We use the conjectured minimax regret beliefs as a starting point for a procedure that selects the strategies more likely to be played by inexperienced agents interacting in one shot 2x2 games. In order to formalize this procedure, we consider a generic 2x2 game where $S_i = \{H_i, T_i\}$ and

$u_i(s_i, s_j)$ are the strategy space and the payoffs of player i with $i, j \in \{A, B\}$. The minimax regret distributions (one for each player) are given by:

$$\{(\tilde{p}_A H_A + (1 - \tilde{p}_A) T_A), (\tilde{p}_B H_B + (1 - \tilde{p}_B) T_B)\}$$

where \tilde{p}_i defines the probability with which player i has to play strategy H_i in order to minimize his regret. With a slightly different notation with respect to the previous sections where only B 's beliefs were considered, define now as $\beta_i = (\theta, 1 - \theta)$ the beliefs of player i such that i believes that player j will play strategies H_j with probability θ and strategy T_j with probability $1 - \theta$. $BR_i(\beta_i)$ is the best reply function of player i . It uses i 's beliefs as an input and provides as an output the strategy that i must choose in order to maximize his expected payoff.

The procedure

1. Compute the minimax regret distribution for the two players and retrieve \tilde{p}_A and \tilde{p}_B .
2. Assign the following beliefs to the two players:

- $\beta_A = (\tilde{p}_B, (1 - \tilde{p}_B))$
- $\beta_B = (\tilde{p}_A, (1 - \tilde{p}_A))$

3. Let the two players choose the strategy to play according to $BR_i(\beta_i)$:

$$\bullet \quad BR_i(\beta_i) = \begin{cases} \{H_i\} & \text{iff } u_i(H_i|\beta_i) > u_i(T_i|\beta_i) \\ \{T_i\} & \text{iff } u_i(H_i|\beta_i) < u_i(T_i|\beta_i) \\ \{0.5H_i + 0.5T_i\} & \text{iff } u_i(H_i|\beta_i) = u_i(T_i|\beta_i) \end{cases}$$

The strategies selected by $BR_i(\beta_i)$ are the ones which have the largest probability to be played in a one shot game or, equivalently, the ones which we would expect to be chosen with the highest frequency if the game is played in a large enough population. In particular whenever $u_i(H_i|\beta_i) \neq u_i(T_i|\beta_i)$ for any i then every player has a single best response and the intersection of the two selected strategies indicates a single outcome of the game as the most likely one to arise.

The procedure thus provides a forecast in three simple steps: it is enough to compute the minimax regret, use its probability distributions to approximate players' beliefs and choose for each player the strategy that best responds to these beliefs.¹² We do not claim this procedure to be consciously used by players. What we claim is that, on average, the procedure is operationally valid, i.e. the majority of individuals play the game "as if" they were applying it.

¹²Notice that the procedure considers all the payoffs of the game. Beliefs of player i are mimicked by the minimax regret probability distribution of the opponent j and thus they depend on the payoffs of the latter. But then, in computing i 's best response, also i 's payoffs are taken into account.

6 Experimental evidences about matching pennies games

We apply the proposed procedure to MP games for which experimental results are available from other studies.¹³ Given that the procedure aims to capture the behavior of inexperienced players, we only consider experiments where subjects played a single game just once. The data are reported in Table 2.

The first three games (*GH1*, *GH2* and *GH3*) and the correspondent experimental results are taken from Goeree and Holt (2001). Each game was played once by a different pool of 50 subjects. In the original paper the authors use these games to evaluate the predictive power of the mixed strategy Nash equilibrium. The last three games appear in Goeree and Holt (2004) who took them from Guyer and Rapoport (1972). In the original experiment 214 subjects played in a random order 244 games belonging to different typologies.

We now explain the meaning of the last four columns of Table 2. In the fourth to last column we report $BR_i(\beta_i)$, the strategy selected by the procedure. The third to last column presents the experimental results in the form $a/b S_i$, where a is the number of players that chose strategy $S_i \in \{H_i, T_i\}$ and $b = 0.5N$ is the total number of row or column players.

The second to last column shows the hit rate which measures the performance of the prediction in forecasting actual behavior. The hit rate is a simple summary statistics which counts the number of hits and it ranges between 0% (all misses) and 100% (all hits).¹⁴ Therefore, when the procedure indicates a single strategy, the hit rate simply captures the percentage of players who actually played it. In games in which the procedure indicates that subjects should uniformly randomize and b is odd, the hit rate reaches 100% if the players split as equally as possible. In game *GH1* for instance the hit rate would have been 100% both if 12 or 13 out of the 25 row or column players chose H_A .

Finally in the last column of Table 2 we test for the validity of our conjecture, i.e. we test the hypothesis of the procedure being able to ex ante predict the strategies that are overplayed. Therefore we only consider the cases in which the procedure selects a pure strategy. Using a one side test, we test if the proportion of players that plays $BR_i(\beta_i)$ is significantly greater than 50%. To do so we use the Fisher's exact probability test which calculates the difference between the data observed and an alternative data distribution. When our procedure selects a pure strategy we expect the null hypothesis (observed data not being significantly different from the uniform distribution) to be rejected. In other words we expect the p-value that appears in the last column to be below 5%. An underlined p-value indicates that this is the case.

¹³With respect to the original papers strategies will be renamed in order to be consistent with previous sections.

¹⁴The hit rate is described in Verbeek (2004) and used for instance in Gneezy and Guth (2003).

Game		Notes		Procedure	Exper.	Hit rate of	Fisher	
N				selects	results	procedure	p-values	
		H_B	T_B					
<i>GH1</i>	H_A	80, 40	40, 80	1 shot	$\frac{1}{2}H_A + \frac{1}{2}T_A$	12/25 H_A	100%	-
	T_A	40, 80	80, 40		$\frac{1}{2}H_B + \frac{1}{2}T_B$	12/25 H_B	100%	-
<i>GH2</i>	H_A	320, 40	40, 80	//	H_A	24/25 H_A	96%	<u>0.04%</u>
	T_A	40, 80	80, 40		T_B	21/25 T_B	84%	<u>1.6%</u>
<i>GH3</i>	H_A	44, 40	40, 80	//	T_A	23/25 T_A	92%	<u>0,18%</u>
	T_A	40, 80	80, 40		H_B	20/25 H_B	80%	<u>3,6%</u>
<i>GR4</i>	H_A	24, 5	5, -10	1 shot	H_A	91/107 H_A	85%	<u>0%</u>
	T_A	26, 9	-10, 26	244 g.	H_B	85/107 H_B	79%	<u>0%</u>
<i>GR5</i>	H_A	15, 5	5, -10	//	H_A	82/107 H_A	77%	<u>0%</u>
	T_A	26, 9	-10, 26		H_B	81/107 H_B	76%	<u>0,01%</u>
<i>GR6</i>	H_A	9, 5	5, -10	//	$\frac{1}{2}H_A + \frac{1}{2}T_A$	74/107 H_A	62%	-
	T_A	26, 9	-10, 26		T_B	32/107 T_B	30%	0,17%

Table 2: the hit rate of the procedure in one shot matching pennies games.

To have a better feeling of how the procedure actually works consider a couple of examples. Game *GH1* is a symmetric matching pennies game. The minimax regret is obviously $\frac{1}{2}H_i + \frac{1}{2}T_i$ for any $i \in \{A, B\}$ and thus the procedure assigns uniform beliefs to both players. Both strategies lead to the same expected payoff and the procedure predicts all outcomes to be equally likely. Even if the population is quite small, actual frequencies confirm that the distributions of players' choices are as uniform as possible.

In game *GH2* the minimax regret distributions are given by $(\frac{7}{8}H_A + \frac{1}{8}T_A)$ and $(\frac{1}{2}H_B + \frac{1}{2}T_B)$. The procedure then selects strategies H_A and T_B given that these are the best responses to the conjectured minimax regret beliefs.¹⁵ Strategy H_A was actually chosen by 24 out of the 25 row players (hit rate of 96%) and strategy T_B was chosen by 21 of the 25 column players (84%).

¹⁵More precisely $H_A \succ T_A$ because $\frac{1}{2}(320) + \frac{1}{2}(40) > \frac{1}{2}(40) + \frac{1}{2}(80)$ and $T_B \succ H_B$ because $\frac{7}{8}(80) + \frac{1}{8}(40) > \frac{7}{8}(40) + \frac{1}{8}(80)$.

The hit rate remains above 75% and the p-values are in line with our conjecture also in games *GH3*, *GR4* and *GR5* while results are less good in the case of Game *GR6* where the procedure failed to predict that *B* players over played strategy H_B .¹⁶ Considering only the games where the procedure indicates a single outcome (*GH2*, *GH3*, *GR4*, *GR5*), the procedure correctly predicts the choices of 81% of the players.

It can be easily shown that predictions based on letting players best respond to beliefs approximated by the *MSNE* clearly under perform our procedure. We already mentioned the fact that using the *MSNE* of player *i* to approximate *j*'s beliefs (Aumann, 1987) does not provide any specific indication given that all the strategies of *j* are a best response to these beliefs. And to use the *MSNE* of player *i* to approximate *i*'s beliefs often leads to misleading results (for instance in Game *GH2* the outcome selected would be (H_A, H_B) instead of (H_A, T_B)).

Moreover also interpreting *MSNE* as a prediction in itself (such that, for instance, players *A* should uniformly randomize in all the three *GH* games given that his mixed equilibrium is always $\frac{1}{2}H_A + \frac{1}{2}T_A$) is not effective. In fact it works fine only in the case of symmetric games (like *GH1*) while it is clearly inadequate in the other games. Indeed Goeree and Holt (2001) present the results of Game *GH1* as supportive of the *MSNE* prediction, while they show the results of games *GH2* and *GH3* as evidences of its failure. Therefore they write that “*The Nash analysis seems to work only by coincidence, when the payoff structure is symmetric and deviation risks are balanced*”.¹⁷

Analyzing the same results through the lens of our conjecture, it seems indeed that the fact that the Nash analysis works in game *GH1* is the result of a coincidence. But this coincidence has an explanation. In symmetric MP games the probability distributions implied by the *MSNE* and by the minimax regret always coincide.¹⁸ Still, as soon as the structure of the game becomes asymmetric, individuals' behavior is by far better captured by our behavioral model rather than by the Nash prediction.

6.1 The procedure in other games

Until now we only considered the more problematic case of MP games but the procedure can be applied to any 2x2 one shot game. The steps to select the strategies that inexperienced players are more likely to choose remain the same: compute the minimax regret, use its probability

¹⁶Note that the payoffs structure of the *GR* games is more complex. Moreover, despite of the fact that games were one shot, the huge number of strategic situations players had to face makes these data less appropriate to study the behavior of inexperienced agents.

¹⁷Goeree, J. & Holt, C. (2001), “Ten Little Treasures of Game Theory and Ten Intuitive Contradictions”, *American Economic Review*, Vol. 91, pp. 1419.

¹⁸See the graphical example in Section 4.4; $k = 1$ identifies the symmetric case, i.e. the unique point for which the functions for the minimax regret and for the *MSNE* intersect.

distribution to approximate players' beliefs and choose the pure strategies that best respond to these beliefs. Table 3 presents examples of a game with a single dominant strategy (*SD*), a prisoner's dilemma (*PD*), a pure coordination game (*PC*), a stag-hunt game (*SH*) and a symmetric (*BS*) and an asymmetric (*aBS*) battle of the sexes.

Game				Minimax regret	Best response	Procedure selects	Notes
		H_B	T_B				
<i>SD</i>	H_A	3, 1	1, 0	$1H_A+0T_A$	H_A	(H_A, H_B)	Unique NE
	T_A	1, 0	0, 2	$\frac{1}{3}H_B+\frac{2}{3}T_B$	H_B		
<i>PD</i>	H_A	3, 3	0, 5	$0H_A+1T_A$	T_A	(T_A, T_B)	Unique NE
		5, 0	1, 1	$0H_B+1T_B$	T_B		
<i>PC</i>	H_A	2, 2	0, 0	$\frac{1}{3}H_A+\frac{2}{3}T_A$	T_A	(T_A, T_B)	Pareto dominant NE
	T_A	0, 0	4, 4	$\frac{1}{3}H_B+\frac{2}{3}T_B$	T_B		
<i>SH</i>	H_A	2, 2	3, 0	$\frac{2}{3}H_A+\frac{1}{3}T_A$	H_A	(H_A, H_B)	Risk dominant NE
	T_A	0, 3	4, 4	$\frac{2}{3}H_B+\frac{1}{3}T_B$	H_B		
<i>BS</i>	H_A	3, 1	0, 0	$\frac{3}{4}H_A+\frac{1}{4}T_A$	$\{H_A, T_A\}$	(\cdot, \cdot)	All outcomes equally likely
	T_A	0, 0	1, 3	$\frac{1}{4}H_B+\frac{3}{4}T_B$	$\{H_B, T_B\}$		
<i>aBS</i>	H_A	5, 1	0, 0	$\frac{5}{6}H_A+\frac{1}{6}T_A$	H_A	(H_A, H_B)	Payoff dominant NE
	T_A	0, 0	1, 3	$\frac{1}{4}H_B+\frac{3}{4}T_B$	H_B		

Table 3: the procedure applied to other classes of 2x2 games.

In accordance with theoretical predictions the procedure selects the unique Nash equilibrium (NE) in the *SD* and *PD* games. For what concerns coordination games the procedure always indicates a single outcome. With this respect it can therefore be considered as a tool for equilibrium selection for games characterized by multiple equilibria. More precisely the procedure selects the Pareto dominant equilibrium in pure coordination games (*PC*). This is in line with intuition, theory and experimental results. More controversial is the indication for stag hunt games (*SH*), i.e. games that have a Pareto dominant NE (more rewarding) and a risk dominant NE (less risky).

The latter is the one indicated by the procedure. For this class of games experimental results provide mixed indications (see for instance Harsanyi and Selten, 1988; Straub, 1995; Haruvy and Stahl, 2004).

More in general, in games that have at least a NE in pure strategies, if the procedure selects a single outcome, then this outcome is always a NE of the game (SD, PD, PC, SH and aBS). However it may be the case that the procedure does not select any outcome (or better it selects them all), even if pure Nash equilibria exist. This is what happens in the case of symmetric battle of the sexes (BS) where the expected payoffs of the two strategies conditional on the conjectured beliefs are equal. The situation is different in the asymmetric version of the game (aBS) where the procedure selects the payoff dominant equilibrium. Both predictions are in line with empirical evidences.

Finally notice that the conjectured minimax regret beliefs sometimes happen to be partially incorrect, i.e. they do not find confirmation in the strategies selected by the opponent (see games SD, BS, aBS). For instance in the SD game the row player expects his opponent to be biased toward playing strategy T_B but indeed, according to the procedure, player B chooses strategy H_B . We do not perceive this to be a problem. In fact we axiomatized the beliefs of inexperienced, unsophisticated and boundedly rational players and therefore the possibility that the procedure may assign incorrect beliefs was embedded in our model since the beginning. What matters is that players actually play as if they were best responding to the minimax regret beliefs such that the procedure is effective in predicting players' behavior. In the case of the SD game for instance, player A chooses his strictly dominant strategy (which is a best response to any possible belief) and player B best responds with strategy H_B .

7 Conclusion

2x2 one shot games remain a fundamental tool for modeling strategic interactions. These games capture the simplest relations (the number of players and strategies is minimal) but still they can be used to describe an uncountable number of situations. No wonder therefore that their study has always attracted a lot of attention. Nevertheless the gap between theoretical models and agents' actual behavior often happens to be still wide. As a consequence, to predict players' behavior in one off interactions remains a problematic issue.

This paper introduced a simple procedure to be used for forecasting the outcome of 2x2 one shot games. Using an axiomatic approach, we looked for a function that may approximate the beliefs of inexperienced and boundedly rational players. First we discussed various proposals connected with the concept of mixed strategy Nash equilibrium and we showed that these functions

cannot be expected to satisfactorily mimic players' beliefs. Then we showed that a belief function based on an instrumental use of minimax regret succeeds in this task. The procedure simply lets players behave as if they were responding to these conjectured minimax regret beliefs. Experimental evidences confirm that the procedure is an effective tool for anticipating the moves of the vast majority of the players.

References

- [1] Aumann, R.J. (1987), "Correlated Equilibrium as an Expression of Bayesian Rationality", *Econometrica*, Vol. 55, No. 1, pp. 1-18.
- [2] Aumann, R.J. and A. Brandenburger (1995), "Epistemic Conditions for Nash Equilibrium", *Econometrica*, Vol. 63, No. 5, pp. 1161-1180.
- [3] Bergemann, D. and K. Schlag (2005), "Robust Monopoly Pricing: the Case of Regret", Cowles Foundation discussion paper no. 1527.
- [4] Borel, E. (1921), "La Théorie du Jeu et les Équations Intégrales à Noyau Symétrique", *Comptes Rendus Hebdomadaires des Séances de l'Académie des Sciences*, Vol. 173, pp. 1304-1308.
- [5] Brafman, R. and M. Tennenholtz (2000), "An Axiomatic Treatment of three Qualitative Decision Criteria", *Journal of the ACM*, Vol. 47, No. 3, pp. 452-482.
- [6] Camerer, C.F. (2003), *Behavioral game theory*, Russell Sage Foundation, NY.
- [7] Camerer, C.F., T. Ho and J.-K. Chong (2004), "A Cognitive Hierarchy Model of Games", *The Quarterly Journal of Economics*, Vol. 119, No. 3, pp. 861-898.
- [8] Chiappori, P.A., S. Levitt and T. Groseclose (2002), "Testing Mixed Strategy Equilibria when Players are Heterogeneous: the Case of Penalty Kicks in Soccer", *American Economic Review*, Vol. 92, No. 4, pp. 1138-1151.
- [9] Costa-Gomes, M., V. Crawford and B. Broseta (2001), "Cognition and behavior in normal form games: an Experimental Study", *Econometrica*, Vol. 69, No. 5, pp. 1193-1235.
- [10] Costa-Gomes, M. and G. Weizsäcker (2005), "Stated Beliefs and Play in Normal Form Games", mimeo, London School of Economics.

- [11] Erev, I. and A. Roth (1998), "Predicting How People Play Games: Reinforcement Learning in Experimental Games with Unique, Mixed Strategies Equilibria", *American Economic Review*, Vol. 88, No. 4, pp. 848-881.
- [12] Gallice, A. (2006a), "Predicting One Shot Play in 2x2 Games using Beliefs based on Minimax Regret", *FEEEM working paper*, No. 31/06.
- [13] Gallice, A. (2006b), "Some Equivalence Results between Mixed Strategy Nash Equilibria and Minimax Regret in 2x2 Games", mimeo, European University Institute.
- [14] Gneezy, U. and W. Guth (2003), "On Competing Rewards Standards - An Experimental Study of Ultimatum Bargaining", *Journal of Socio-Economics*, Vol. 31, No. 6, pp. 599-607.
- [15] Goeree, J.K. and C. Holt (2001), "Ten Little Treasures of Game Theory and Ten Intuitive Contradictions", *American Economic Review*, Vol. 91, No. 5, pp. 1402-1422.
- [16] Goeree, J.K., C. Holt and T. Palfrey (2003), "Risk Averse Behavior in Generalized Matching Pennies", *Games and Economic Behavior*, Vol. 45, No. 1, pp. 97-113.
- [17] Goeree, J.K. and C. Holt (2004), "A Model of Noisy Introspection", *Games and Economic Behavior*, Vol. 46, No. 2, pp. 365-382.
- [18] Guyer, M. and A. Rapoport (1972), "2x2 Games Played Once", *The Journal of Conflict Resolution*, Vol. 16, No. 3, pp. 409-431.
- [19] Haruvy, E. and D.O. Stahl (2004), "Deductive versus Inductive Equilibrium Selection: Experimental Results", *Journal of Economic Behavior and Organization*, Vol. 53, No. 3, pp. 319-331.
- [20] Harsanyi, J. and R. Selten (1988), *A General Theory of Equilibrium Selection in Games*, M.I.T. Press.
- [21] Manski, C. (2005), "Minimax-Regret Treatment Choice with Missing Outcome Data", *Journal of Econometrics*, forthcoming.
- [22] McKelvey, R. and T. Palfrey (1995), "Quantal Response Equilibria for Normal Form Games", *Games and Economic Behavior*, Vol. 10, No. 1, pp. 6-38.
- [23] McKelvey, R., T. Palfrey and R. Weber (2000), "The Effects of Payoff Magnitude and Heterogeneity on Behavior in 2x2 Games with Unique Mixed Strategy Equilibria", *Journal of Economic Behavior and Organization*, Vol. 42, No. 4, pp. 523-548.

- [24] Milnor, J. (1954), "Games against Nature", Ch. IV, pp. 49-60, in R.M. Thrall, C.H. Coombs, and R.L. Davis (eds) *Decision Processes*, Wiley and Chapman & Hall, New York and London.
- [25] Mookherjee, D. and B. Sopher (1994), "Learning Behavior in an Experimental Matching Pennies Game", *Games and Economic Behavior*, Vol. 7, No. 1, pp. 62-91.
- [26] Nash, J. (1951), "Non-Cooperative Games", *Annals of Mathematics*, Vol. 54, No. 2, pp. 286-295.
- [27] Nyarko, Y. and A. Schotter (2002), "An Experimental Study of Belief Learning using Elicited Beliefs", *Econometrica*, Vol. 70, No. 3, pp. 971-1005.
- [28] Ochs, J. (1995), "Games with Unique, Mixed Strategy Equilibria: An Experimental Study", *Games and Economic Behavior*, Vol. 10, No. 1, pp. 202-217.
- [29] Osborne, M.J. and A. Rubinstein (1994), *A course in game theory*, M.I.T. Press.
- [30] Palacios-Huerta, I. (2003), "Professionals Play Minimax", *Review of Economic Studies*, Vol. 70, No. 2, pp. 395-415.
- [31] Reny, P.J. and A.J. Robson (2004), "Reinterpreting Mixed Strategy Equilibria: A Unification of the Classical and Bayesian Views", *Games and Economic Behavior*, Vol. 48, No. 2, pp. 355-384.
- [32] Rey-Biel, P. (2004), "Equilibrium Play and Best Response to (Stated) Beliefs in Constant Sum Games", mimeo, University College London.
- [33] Rosenthal, R. W. (1979), "Sequences of Games with Varying Opponents", *Econometrica*, Vol. 47, No. 6, pp. 1353-1366.
- [34] Roth, A.E. (1995), "Introduction to Experimental Economics", Ch I, pp. 3-109, in J.H. Kagel and A.E. Roth (eds), *The Handbook of Experimental Economics*, Princeton University Press, Princeton.
- [35] Rubinstein, A. (1991), "Comments on the Interpretation of Game Theory", *Econometrica*, Vol. 59, No. 4, pp. 909-924.
- [36] Savage, L. (1951), "The Theory of Statistical Decision", *Journal of the American Statistical Association*, Vol. 46, No. 253, pp. 55-67.
- [37] Savage, L. (1954), *The Foundations of Statistics*, John Wiley and Sons, New York.

- [38] Stahl, D.O. and P.W. Wilson (1995), “On Players Model of Other Players: Theory and Experimental Evidence”, *Games and Economic Behavior*, Vol. 10, No. 1, pp. 218-254.
- [39] Straub, R. (1995), “Risk Dominance and Coordination Failures in Static Games”, *The Quarterly Review of Economics and Finance*, Vol. 35, No. 4, pp. 339-363.
- [40] Tang, F.-F. (2001), “Anticipatory Learning in two-persons Games: some Experimental Results”, *Journal of Economic Behavior and Organization*, Vol. 44, No. 2, pp. 221-232.
- [41] Verbeek, M. (2004), *A guide to Modern Econometrics*, John Wiley and Sons, New York.
- [42] Von Neumann, J. and O. Morgenstern (1944), *Theory of Games and Economic Behavior*, Princeton University Press, Princeton.
- [43] Walker, M. and J. Wooders (2001), “Minimax Play at Wimbledon”, *American Economic Review*, Vol. 91, No. 5, pp. 1521-1528.