Minmax regret and deviations from Nash Equilibrium

Michal Lewandowski*

December 10, 2019

Abstract

We build upon Goeree and Holt [*American Economic Review*, 91 (5) (2001), 1402-1422] and show that the departures from Nash Equilibrium predictions observed in their experiment on static games of complete information can be explained by minimizing the maximum regret.

Keywords: minmax regret, Nash equilibrium
JEL Classification Numbers: C7

1 Introduction

1.1 Motivating examples

Game theory predictions are sometimes counter-intuitive. Below, we give three examples. First, consider the traveler's dilemma game of Basu (1994). An airline loses two identical suitcases that belong to two different travelers. The airline worker talks to the travelers separately asking them to report the value of their case between \$2 and \$100. If both tell the same amount, each gets this amount. If one amount is smaller, then each of them will get this amount with either a bonus or a malus: the traveler who chose the smaller amount will get \$2 extra; the other traveler will have to pay \$2 penalty.

Intuitively, reporting a value that is a little bit below \$100 seems a good strategy in this game. This is so, because reporting high value gives you a chance of receiving large payoff without risking much – compared to reporting lower values the maximum possible loss is \$4, i.e. the difference between getting the bonus and getting the malus. Bidding a little below instead of exactly \$100

*Warsaw School of Economics, michal.lewandowski@sgh.waw.pl

Figure 1: Asymmetric matching pennies

	L	R
T	$\underline{7}, 0$	$0, \underline{1}$
B	$0, \underline{1}$	$\underline{1}, 0$

Figure 2: Choose an effort games

Game A

	L	R
T	$\underline{2}, \underline{2}$	$-3, 1$
B	$1, -3$	$\underline{1}, \underline{1}$

Game B

	L	R
T	$\underline{5}, \underline{5}$	$0, 1$
B	$1, 0$	$\underline{1}, \underline{1}$

avoids choosing a strictly dominant action. In striking contrast to that intuition, the only profile of rationalizable actions is for both players to report the minimal value of \$2. Note that by playing a rationalizable action you can never receive more than \$4 payoff, whereas playing a little bit below \$100 gives you a chance of getting \$100.

Consider the second example, an asymmetric matching pennies game whose payoff table is given in Figure 1. The only Nash equilibrium of this game is in mixed strategies: $\left(\frac{1}{2}T + \frac{1}{2}B, \frac{1}{8}L + \frac{7}{8}R\right)$. Note that as compared to the symmetric matching pennies game (replace 7 with 1 in the above game), it is only the column player who "reacts" to this asymmetry of payoffs by increasing the probability of playing R. The row player did not adjust his strategy, even though it is his payoffs that became asymmetric, and not those of his opponent. This property of a mixed strategy equilibrium, in which the player's strategy is calculated based on the other player's (and not his own) payoff also seems counter-intuitive as compared to what we would expect people to play in this game.

As the third example, consider two games of a "choose-an-effort" variety whose payoff tables are given in Figure 2. In both these games there are: two pure strategy equilibria (T, L) and (B, R), and one mixed strategy equilibrium: $(\frac{4}{5}T + \frac{1}{5}B, \frac{4}{5}L + \frac{1}{5}R)$ in game A and $(\frac{1}{5}T + \frac{4}{5}B, \frac{1}{5}L + \frac{4}{5}R)$ in game B. Game theory does not provide an answer which of the three equilibria will be played. Except that the equilibrium (T, L) is Pareto efficient in both of these games and so maybe it will be played more often. Yet, when trying to predict the play between real people, we may argue as follows. In game A, playing T is risky as compared to playing B (payoff is either 2 or -3 in the former case and 1 in the latter). The same holds for a column player: playing L is risky as compared to playing R. So, intuitively, we expect more people to play (B, R) than (T, L). On the other hand, in game B, playing T is not that risky as compared to playing B (5 or 0 in the former case as compared to 1 in the latter). The same for the column player: playing L is not that risky compared to R. Intuitively, we expect more people to play (T, L) than (B, R) in this game.

Note that this intuition is in striking contrast with the mixed strategy equilibrium which predicts exactly the opposite. It advises players to play T more often than B in game A and B more often than T in game B.

1.2 Goal of the paper: the minmax regret hypothesis

The implicit assumption underlying the Nash Equilibrium (and rationalizability) is that of common knowledge among the players of the rules of a game, preferences of the players and the rationality of the players. Rationality of a player is defined as consistency in pursuing his objective whereas the objective is assumed to be the maximization of the expected value of the player's own payoff, measured on some utility scale (Myerson, 1991, p. 2). Common knowledge of a fact p among the players (Aumann, 1976) requires that every statement of the form: "(every player know that)^k every player knows p" is true, for $k = 0, 1, 2, ...$

In experiments testing game theory, rules of a game may be made common knowledge. However, rationality of the players and their exact motives (preferences) is usually not a common knowledge. Players usually do not know each other, they might be making mistakes, lose their focus or interest for a while, be uncertain about other players' motives, etc. Since Nash equilibrium is heavily based on the assumption of common knowledge of rationality among the players, we might observe departures from Nash equilibrium predictions. And if so, what should we expect people to play and why? Even if players adhere to a Nash equilibrium strategy, there are games with many equilibrium strategies. What makes people choose one equilibrium strategy over another. Are they all equally attractive? Or maybe there is something that makes one of them more attractive or intuitive over another?

By departing from the assumption of common knowledge of rationality and of other players' preferences, an additional uncertainty arises - players become less predictable in their behavior. In the extreme case when his opponent is believed to be completely unpredictable, a player might treat a strategic two-player game as a game against nature under conditions of complete ignorance. In such a case the player might disregard his opponent's payoffs, treat his

Figure 3: Regret tables for choose an effort games

Game A

	L	R	max
T	0, 0	4, 1	4
B	1, 4	0, 0	**1**
max	4	**1**	

Game B

	L	R	max
T	0, 0	1, 4	**1**
B	4, 1	0, 0	4
max	**1**	4	

strategies as states of nature and choose according to one of the rules of choice under complete ignorance. We argue that, even though it is the extreme case, it may help predicting how people actually play in experimental settings.

We postulate using the minmax regret decision rule of Savage (1951) to help predicting departures from Nash Equilibrium. Among many decision rules of choice under complete ignorance we chose the minmax regret rule for a number of reasons. First, unlike the Hurwicz criterion (Hurwicz, 1951) with the pessimism coefficient α, it does not require eliciting any preference parameters. Second, it is neither overly pessimistic as the Wald's maxmin rule (Wald, 1950) nor overly optimistic as the maxmax rule. It also avoids the problem of the Laplace's principle of insufficient reason (Laplace, 1814) criticized for not reflecting properly the idea of complete ignorance (uniform prior is not uninformative). Finally, in the context of a game the minmax regret rule does not exhibit its main drawback. The rule generally violates independence of irrelevant alternatives (i.e. presence of an unwanted alternative might change the choice), but in a game the choice alternatives are fixed, so the issue of choice reversals due to adding or removing choice alternatives to the choice set does not arise.

We now show that the minmax regret rule captures the intuition behind the selection of strategies that we presented in the introduction. Consider games A and B. Their regret tables are given in Figure 3. The minmax regret strategies for both players are: B and R in the case of game A and T and L in the case of game B. This is consistent with the intuition that was given in the introduction. In general, we will argue that minmax regret captures well the intuition provided while discussing the motivating examples.

The hypothesis we test is whether minmax regret and the intuition it captures helps predicting the actual play in simple static games of complete information. In particular, we test whether players move towards strategies with low minmax regret if the difference in regrets between strategies is large enough. For this purpose we analyze experimental data obtained by Goeree and Holt

(2001). They report data for a series of two-person games played once. For each game they assume two payoff structures: the treasure treatment in which the observed behavior agrees with the NE prediction, and the contradiction treatment in which the behavior strikingly deviates from it. They analyze static and dynamic games with complete and incomplete information. Here we restrict attention only on their static games of complete information.

1.3 Related literature

Using minmax regret rule in the context of a game was originally proposed by Luce and Raiffa as early as in 1957 in their classic book Games and decisions. They state that in the context of a strategic interaction independence of irrelevant alternatives can be criticized because "adding a new act for the decision maker can affect the strategic position of the adversary and therefore the decision maker should reappraise the relative merits of the old facts". Hence they postulate that the minmax regret criterion, "which was mainly criticized on the basis of its non-independence of irrelevant alternatives, should be reevaluated" (Luce and Raiffa, 1989, p. 307). Renou and Schlag (2010) take up this challenge and propose the minmax regret equilibrium notion. While we share their motivation that minmax regret captures the intuition of what happens when we drop the assumption of mutual common knowledge of rationality, our goal in this paper is different than theirs. They propose the equilibrium notion, we suggest to use the rule to help predicting departures from Nash Equilibria. Our use of the rule is non-strategic whereas they treat it as part of strategic considerations.

2 Formal definitions and notation

Consider a game Γ in strategic form:

1. The set of players N
2. For each player $i \in N$ a set of actions A_i
3. For each player $i \in N$ a payoff function $u_i : A \to \mathrm{R}$

Notation:

- A set of all players profiles of actions is denoted by $A \equiv \times_{i\in N} A_i$ with a typical element denoted by a.

- A set of all players but player i profiles of actions is denoted by $A_{-i} \equiv \times_{j \in N \setminus i}$ with a typical element denoted by a_{-i}.

Definition 1. *Nash equilibrium in pure strategies is a profile of actions $a^* \in A$ such that:*

$$a_i^* \in \arg\max_{a_i \in A_i} u_i(a_i, a_{-i}^*), \quad \forall a_{-i} \in A_{-i}, \ \forall i \in N \tag{1}$$

Definition 2. *Given any strategic form game Γ, a randomized strategy for any player i is a probability distribution over A_i. Let $\Delta(A_i)$ denote the set of all possible randomized strategies for player i. The set of all randomized strategy profiles will be denoted by $\Delta(A) = \times_{i \in N} \Delta(A_i)$. It must be that:*

$$\sum_{a_i \in A_i} \sigma_i(a_i) = 1, \quad \forall i \in N$$

We will write $\sigma \equiv (\sigma_i)_{i \in N}$, where $\sigma_i \equiv (\sigma_i(a_i))_{a_i \in A_i}$, for each i.
For any randomized strategy profile σ, let $u_i(\sigma)$ denote the expected payoff that player i would get when the players independently choose their pure strategies according to σ:

$$u_i(\sigma) = \sum_{a \in A} \left(\prod_{j \in N} \sigma_j(a_j) \right) u_i(a), \quad \forall i \in N$$

For any $\sigma_i' \in \Delta(A_i)$, we denote (σ_i', σ_{-i}) the randomized strategy profile in which the i-th component is σ_i' and all other components are as in σ. Thus:

$$u_i(\sigma_i', \sigma_{-i}) = \sum_{a \in A} \left(\prod_{j \in N \setminus i} \sigma_j(a_j) \right) \sigma_i'(a_i) u_i(a)$$

Definition 3. *A randomized strategy profile $\sigma^* \in \Delta(A)$ is a Nash equilibrium of Γ if the following holds:*

$$\sigma_i^* \in \arg\max_{\sigma_i \in \Delta(A_i)} u_i(\sigma_i, \sigma_{-i}^*), \quad \forall i \in N \tag{2}$$

Definition 4. *The rationalizable set of actions (Bernheim, 1986, Pearce, 1984) can be computed as follows:*

1. *Start with the full action set for each player.*

2. *Remove all actions which are never a best reply to any belief about the opponents' actions – no rational player will choose such actions.*

3. *Remove all actions which are never a best reply to any belief about the opponents' remaining actions – no player who knows that the other players are rational will choose such actions.*

4. *Continue the process until no further actions are eliminated.*

5. *In a game with finitely many actions, this process always terminates and leaves a non-empty set of actions for each player.*

Players respond optimally to some belief about their opponents' actions, but Nash equilibrium requires that these beliefs be correct while rationalizability does not. The general idea is to provide the weakest constraints on players while still requiring that players are rational and this rationality is common knowledge among the players. We now define regret of a strategy

Definition 5. *Regret of an action* $a_i' \in A_i$ *given a profile of other players' actions* $a_{-i}' \in A_{-i}$ *is defined as:*

$$R_i(a_i', a_{-i}') = \max_{a_i \in A_i}[u_i(a_i, a_{-i}')] - u_i(a_i', a_{-i}') \tag{3}$$

The maximum regret of $a_i' \in A_i$ *is then given by:*

$$\max_{a_{-i} \in A_{-i}} R_i(a_i', a_{-i}) \tag{4}$$

The minmax regret strategy for a player $i \in N$ *is an action* $a_i^* \in A_i$ *such that:*

$$a_i^* \in \arg\min_{a_i \in A_i} \left[\max_{a_{-i} \in A_{-i}} R_i(a_i, a_{-i}) \right] \tag{5}$$

3 Minmax regret strategies and treasures of game theory

In this section we analyze those of the games from Goeree and Holt (2001) that are static games of complete information.

3.1 Matching pennies

Consider three games of a "matching pennies" variety – their payoff and regret tables are given in Figure 4. The first game is symmetric whereas the other two are asymmetric. Note that, as discussed in the introduction, since the column player's payoffs are the same in all three games, the row player's equilibrium strategy is the same in all three games. It sharply contrasts with the experimental evidence reported by Goeree and Holt (2001). In the asymmetric matching

Figure 4: Matching pennies. In brackets: mixed Nash equilibrium in BLACK, experimental data in RED.

Symmetric matching pennies

	L (.50)/(.48)	R
T (.50)/(.48)	<u>80</u>, 40	40, <u>80</u>
B	40, <u>80</u>	<u>80</u>, 40

Regret table

	L	R	max
T	0, 40	40, 0	40
B	40, 0	0, 40	40
max	40	40	

Asymmetric matching pennies

	L (.13)/(.16)	R
T (.50)/(.96)	<u>320</u>, 40	40, <u>80</u>
B	40, <u>80</u>	<u>80</u>, 40

Regret table

	L	R	max
T	0, 40	40, 0	**40**
B	280, 0	0, 40	280
max	40	40	

Reversed asymmetry

	L (.91)/(.80)	R
T (.50)/(.08)	<u>44</u>, 40	40, <u>80</u>
B	40, <u>80</u>	<u>80</u>, 40

Regret table

	L	R	max
T	0, 40	40, 0	40
B	4, 0	0, 40	**4**
max	40	40	

pennies game in the middle panel people tend to choose strategy T twenty four times more often than strategy B. In the reversed asymmetry game in the bottom panel it is exactly the opposite: strategy B is played more than twelve times more often than strategy T.

The data pattern is captured by the minmax regret rule. In the “asymmetric matching pennies” game the maximum regret of strategy T is seven times lower than that of strategy B whereas in the “reversed asymmetry” game it is the opposite: the maximum regret of strategy T is ten times higher than that of strategy B. Note also that as for the column player strategy, experimental data is roughly consistent with the equilibrium prediction. It is also the case that for the column player the maximum regret values of the two strategies are the same.

3.2 The traveller's dilemma

Let us consider the travellers' dilemma with two players and action space $A_i = \{180, 181, ..., 300\}$ for each $i \in N = \{1, 2\}$. Payoffs are the following:

$$u_i(a_i', a_j') = \min(a_i', a_j') + P(a_i', a_j'), \quad i \neq j, \quad i, j \in N,$$

$$\text{where} \quad P(a_i', a_j') = \begin{cases} P & \text{if} \quad a_i' < a_j' \\ 0 & \text{if} \quad a_i' = a_j' \\ -P & \text{if} \quad a_i' > a_j' \end{cases}, \quad \text{where} \quad P \in \mathrm{Z}^+$$

When $P = 0$ and for $P = 1$, any profile of strategies for which $a_i = a_j, \ \ i \neq j$ is a pure strategy Nash equilibrium. When $P > 1$, the only Nash equilibrium is for every player to bid 180. It is also the unique profile of rationalizable actions in this game. Experimental evidence shows that if P is equal to 180 almost 80% of all players report values in the lowest interval of $180 - 190$. However, if P is equal to 5, then almost 80% of all players report values in the highest interval of $290 - 300$.

We now calculate the max regret of all the strategies in this game. Regret of player i for a given strategy profile $(a_i', a_j') \in A$ is equal to:

$$R_i(a_i', a_2') = \max_{a_i \in A_i} \big(\min(a_i, a_j') + P(a_i, a_j')\big) - \big(\min(a_i', a_j') + P(a_i', a_j')\big)$$

We need to consider two cases:

$$a_j' = 180 \ \Rightarrow \ R_i(a_i', 180) = \begin{cases} 180 - 180 = 0 & \text{if} \quad a_i' = 180 \\ 180 - 180 + P = P & \text{if} \quad a_i' > 180 \end{cases}$$

$$a_j' > 180 \ \Rightarrow \ R_i(a_i', a_j') = a_j' - 1 + P - \big(\min(a_i', a_j') + P(a_i', a_j')\big)$$

Let's summarize it in the form of the table:

	$a_i' = 180$	$a_i' > 180$
$a_j' = 180$	0	P
$a_j' > 180$	$a_j' - 181$	$a_j' - 1 + P - \min(a_i', a_j') - P(a_i', a_j')$

So the maximum regret of player i for a given strategy $a_i' \in A_i$ is equal to:

$$\max \text{regret}_i(a_i') = \begin{cases} 119 & \text{when} \quad a_i' = 180 \\ \max(P, 118) & \text{when} \quad a_i' = 181 \\ \max(2P - 1, 0) & \text{when} \quad a_i' \geq 182 \end{cases}$$

Now we can solve for the minmax strategies in this game. They are summarized in Table 1. For example, for $P = 0$, the only minmax regret strategy is to bid 300. For $P = 5$, any bid in the set $\{290, 291, ..., 300\}$ is a minmax regret strategy. For $P = 180$ the only minmax regret strategy is to bid 180.

Table 1: Minmax regret strategies of the traveller's dilemma as a function of P

Values of P	Set of minmax regret strategies	minmax value
$P = 0$	$\{300\}$	0
$P \in \{1, ..., 59\}$	$\{300 - 2P, ..., 299, 300\}$	$2P - 1$
$P \in \{60, 61, ..., 118\}$	$\{181\}$	118
$P = 119$	$\{180, 181\}$	119
$P \in \{120, 121, ...\}$	$\{180\}$	119

3.3 Choose an effort game

The strategy space is $A_i = \{110, 111, ..., 170\}, \ i \in N = \{1, 2\}$. The payoffs for a given profile of actions $(a'_i, a'_j) \in A$ are a function of a cost of effort parameter $c \in (0, 1)$:

$$u_i(a'_i, a'_j) = \min(a'_i, a'_j) - ca'_i$$

The set of Nash equilibria consists of all the pairs (a, a), where $a \in A_i$. The experimental evidence on the other hand suggests that if cost of effort is low ($c = 0.1$), then people usually choose high effort whereas if cost of effort is high ($c = 0.9$), then they choose low effort. Again this is not captured by the NE prediction. Let's calculate the riskiness of the strategies involved. The regret of a given profile of actions $(a'_i, a'_j) \in A$ is given by:

$$\begin{aligned} R_i(a'_i, a'_j) &= \max_{a_i \in A_i} \big(\min(a_i, a'_j) - ca_i\big) - \big(\min(a'_i, a'_j) - ca'_i\big) \\ &= a'_j - ca'_j - \min(a'_i, a'_j) + ca'_i \end{aligned}$$

The maximum regret of a strategy a'_i is given by:

$$\begin{aligned} \max \text{regret}_i(a'_i) &= \max_{a_j \in A_j} \big(a_j - ca_j - \min(a'_i, a_j) + ca'_i\big) \\ &= \max\big(170(1 - c) - a'_i, -110c\big) + ca'_i \end{aligned}$$

The minmax regret is then:

$$\min_{a_i \in A_i} [\max (170(1 - c) - a_i, -110c) + ca_i]$$

Let's define $a^*_i \in A_i$ as the value of player i strategy for which the two elements of the above max function are equal:

$$170(1 - c) - a^*_i = -110c \quad \Rightarrow \quad a^*_i = 170 - c(170 - 110)$$

Figure 5: A coordination game with a secure outside option. In brackets: mixed NE in BLACK, experimental data in RED

$x = \mathbf{0}$	L (.67)	M (.33)/(.84)	R
T	$\underline{90}, \underline{90}$	$0, 0$	$\mathbf{0}, 40$
B (.33)/(.96)	$0, 0$	$\underline{180}, \underline{180}$	$0, 40$

$x = \mathbf{400}$	L (.67)	M (.33)/(.76)	R
T	$\underline{90}, \underline{90}$	$0, 0$	$\mathbf{400}, 40$
B (.33)/(.64)	$0, 0$	$\underline{180}, \underline{180}$	$0, 40$

Regret table	L	M	R	max
T	$0, 0$	$180, 90$	$0, 50$	180
B	$90, 180$	$0, 0$	$x, 140$	$\max(90, x)$
max	180	**90**	140	

In order to find the minmax regret, it is clear that we need to consider three cases:

$$\begin{aligned} a_i' = 110 \quad &\Rightarrow \quad \max \text{regret}_i(110) = (1-c)(170-110) \\ a_i' = a_i^* \quad &\Rightarrow \quad \max \text{regret}_i(a_i^*) = c(1-c)(170-110) \\ a_i' = 170 \quad &\Rightarrow \quad \max \text{regret}_i(170) = c(170-110) \end{aligned}$$

Since $c \in (0,1)$ it must be that the minmax regret strategy is $a_i' = a_i^*$ and the minmax regret value is $\max \text{regret}_i(a_i^*) = c(1-c)(170-110)$. For example if $c = 0.1$, the minmax regret strategy is equal to 164 and if $c = 0.9$, the minmax regret strategy is equal to 116, which is in accordance with the experimental evidence.

3.4 The extended coordination game

Figure 3.4 presents the payoffs, data as well as the minmax regret strategies of the extended coordination game. Note that according to the evidence people most frequently played (B, M) profile of strategies, i.e. the Pareto efficient (pure strategy) Nash equilibrium. However, row players decreased their frequency of playing B by 32% in the 400 treatment as compared to the 0 treatment. This change is predicted by the minmax regret rule: B is the minmax regret strategy in the 0 treatment and T becomes the minmax regret strategy in the 400 treatment.

Figure 6: The Kreps game. In brackets: mixed NE in black, experimental data in red.

	L (.13)/(.26)	M (.08)	NN (.68)	R (.87)/(.00)
T (.49)/(.68)	$\underline{200}, \underline{50}$	$0, 45$	$10, 30$	$20, -250$
B (.51)/(.32)	$0, -250$	$\underline{10}, -100$	$\underline{30}, 30$	$\underline{50}, \underline{40}$

Regret table	L	M	NN	R	max
T	$0, 0$	$10, 5$	$20, 20$	$30, 300$	**30**
B	$200, 290$	$0, 140$	$0, 10$	$0, 0$	200
max	290	140	**20**	300	

3.5 The Kreps game

Figure 3.5 presents the payoffs, data and the minmax regret strategies in the so called Kreps game. Note that the strategy NN of the Column player is not a best-response to any strategy of the Row player. Yet, it is played most frequently. This may be explained by the fact, that all other strategies of the column player are risky - their maximum regret ranges between $140 - 300$. This is not so for the strategy NN for which the max regret is minimal and equals 20.

4 Conclusion

We have analyzed static games of complete information from Goeree and Holt (2001) and found that the deviations from Nash Equilibrium strategies observed in their data may be explained by taking into account the strategies' maximum regret. The rationale behind this criterion stems from relaxing the assumption of common knowledge of rationality. This introduces the uncertainty about the opponents behavior that is well captured by the maximum regret.

Maximum regret should be seen as an auxiliary criterion based on which people decide which strategy to choose, the first being the Nash equilibrium strategy. It should not be treated as a binary criterion - the magnitude of regret should also matter.

Future research should also analyze the second order minmax regret strategy, i.e. the best response to the minmax regret strategy of the opponent.

References

Aumann, R. J. (1976). Agreeing to disagree. *The annals of statistics*, 1236–1239.

Basu, K. (1994). The traveler's dilemma: Paradoxes of rationality in game theory. *The American Economic Review 84*(2), 391–395.

Bernheim, B. D. (1986). Axiomatic characterizations of rational choice in strategic environments. *The scandinavian journal of economics*, 473–488.

Goeree, J. K. and C. A. Holt (2001). Ten little treasures of game theory and ten intuitive contradictions. *The American Economic Review 91*(5), 1402–1422.

Hurwicz, L. (1951). The generalized bayes minimax principle: a criterion for decision making under uncertainty. *Cowles Comm. Discuss. Paper Stat 335*, 1950.

Laplace, P. (1814). A philosophical essay on probabilities.

Luce, R. D. and H. Raiffa (1989). *Games and decisions: Introduction and critical survey*. Dover publications: Reprint.

Pearce, D. G. (1984). Rationalizable strategic behavior and the problem of perfection. *Econometrica: Journal of the Econometric Society*, 1029–1050.

Renou, L. and K. H. Schlag (2010). Minimax regret and strategic uncertainty. *Journal of Economic Theory 145*(1), 264–286.

Savage, L. J. (1951). The theory of statistical decision. *Journal of the American Statistical association 46*(253), 55–67.

Stoye, J. (2011). Axioms for minimax regret choice correspondences. *Journal of Economic Theory 146*(6), 2226–2251.

Wald, A. (1950). Statistical decision functions.