University *of* Leicester

Implementation in Minimax Regret Equilibrium

Ludovic Renou, University of Leicester, UK

Karl H. Schlag, Universitat Pompeu Fabra

Implementation in Minimax Regret Equilibrium[*]

Ludovic Renou [†] & Karl H. Schlag [‡]

October 12, 2009

Abstract

This note studies the problem of implementing social choice correspondences in environments where individuals have doubts about the rationality of their opponents. We postulate the concept of ε-minimax regret as our solution concept and show that social choice correspondences that are Maskin monotonic and satisfy the no-veto power condition are implementable in ε-minimax regret equilibrium for all $\varepsilon \in [0, 1)$.

Keywords: implementation, minimax regret, Maskin monotonicity.

JEL Classification Numbers: C72; D71

[†]Department of Economics, Astley Clarke Building, University of Leicester, University Road, Leicester LE1 7RH, United Kingdom. `lr78@le.ac.uk`

[‡]Department of Economics and Business, Universitat Pompeu Fabra, Spain`karl.schlag@upf.edu`

1 Introduction

This note studies the problem of implementing social choice correspondences in environments where individuals have doubts about the rationality of their opponents.

Now, if an individual is uncertain about the rationality of his opponents, which conjectures about his opponents' actions should he form? This is a very intricate issue. Admittedly, he can form a subjective probabilistic assessment and play a best response to his assessment. However, any subjective assessment is largely arbitrary, and there is no obvious reasons to favor one assessment over another. Bayesian theory is silent on how to form initial probabilistic assessments. Moreover, experimental evidence such as the Ellsberg paradox suggests that individuals frequently experience difficulties in forming a unique assessment. In this note, we postulate that "regret" guides individuals in forming probabilistic assessments and, ultimately, in making choices. Specifically, we consider the concept of ε-minimax regret equilibrium as the solution concept (Renou and Schlag (2009)).

In an ε-minimax regret equilibrium, each player believes that his opponents are playing according to the equilibrium strategies with probability $1 - \varepsilon$, and is completely uncertain about the play of his opponents, otherwise. Whenever uncertain about the play of his opponents, a player conjectures that his opponents would play so as to maximize his regret. In particular, there is no mutual belief in rationality in an ε-minimax regret equilibrium (unless $\varepsilon = 0$). Intuitively, if a player believes with probability one that his opponents are rational, he must conclude that his opponents will not play strictly dominated strategies. However, this might contradict his conjecture that his opponents aims at maximizing his regret with probability ε: strictly dominated strategies might maximize his regret. The parameter ε thus captures the extent to which players are doubtful about the rationality of others: the higher ε is, the more doubtful a player is.

We show that the social choice correspondences that are Maskin monotonic and satisfy the no-veto power condition are implementable in ε-minimax regret equilibrium for any $\varepsilon < 1$. Perhaps surprising, this result states that even arbitrarily large uncertainties about the rationality of others does not undermine the implementation of social correspondences that are Maskin monotonic and satisfy no-veto power. We also show that Maskin monotonicity is not a necessary condition for implementation in ε-minimax regret equilibrium, even when implementation is required for all $\varepsilon \in (0, 1)$. A larger set

of social choice correspondences can be implemented.

For excellent surveys on implementation theory, we refer the reader to Jackson (2001) and Maskin and Sjöström (2002). A closely related contribution to our work is Tumennasan (2008), who studies the problem of implementation in quantal response equilibrium. The distinctive feature of this solution concept is that players make errors in evaluating their payoffs and, thus, may play sub-optimal strategies. Tumennasan makes, however, two additional and key assumptions. Firstly, he considers limiting quantal response equilibrium, i.e., as the probability of errors goes to zero. Secondly, he requires each limiting quantal response equilibrium to be a strict Nash equilibrium. These two additional assumptions makes the problem of implementation in quantal response equilibrium essentially equivalent to the problem of implementation in strict Nash equilibria. He shows that quasi-monotonicity together with a condition termed no-worst alternatives are necessary and almost sufficient conditions for implementation.[1] Quasi-monotonicity and Maskin monotonicity do not imply each others. Unlike Tumennasan, we do not refine the set of ε-minimax regret equilibria and, in fact, show that neither Maskin monotonicity nor quasi-monotonicity are necessary conditions for implementation in ε-minimax regret equilibrium.

2   Preliminaries

An environment is a triplet $\langle N, X, \Theta \rangle$ where $N := \{1, \ldots, n\}$ is a set of n players, X a finite set of alternatives, and Θ a finite set of states of the world. For each player $i \in N$, there is a Von Neumann-Morgenstern utility function $u_i : X \times \Theta \to \mathbb{R}$.

We denote by $L_i(x, \theta) := \{y \in X : u_i(x, \theta) \geq u_i(y, \theta)\}$ player i's lower contour set of x at state θ.

A social choice correspondence $f : \Theta \to 2^X \setminus \{\emptyset\}$ associates with each state of the world θ a non-empty subset of alternatives $f(\theta) \subseteq X$. Two classic conditions for Nash implementation are Maskin monotonicity and no-veto power. A social choice correspondence f is *Maskin monotonic* if for all (x, θ, θ') in $X \times \Theta \times \Theta$ with $x \in f(\theta)$, we have $x \in f(\theta')$ whenever $L_i(x, \theta) \subseteq L_i(x, \theta')$ for all $i \in N$. Maskin monotonicity is a

---

[1]For sufficiency, he also needs the no-veto power condition. For the problem of implementation in strict Nash equilibrium, Cabrales and Serrano (2008) show that quasi-monotonicity and the condition of no-worst alternatives are necessary and almost sufficient conditions.

necessary condition for Nash implementation. A social choice correspondence f satisfies *no-veto power* if for all $\theta \in \Theta$, we have $x \in f(\theta)$ whenever $x \in \arg\max_{x' \in X} u_i(x', \theta)$ for all but at most one player $i \in N$. Maskin monotonicity and no-veto power are sufficient conditions for Nash implementation (if $n \geq 3$). Denote $\mathcal{F}_{NE}$ the set of social choice correspondences that are Maskin monotonic and satisfies the no-veto power condition.

A mechanism (or game form) is a pair $\langle (M_i)_{i \in N}, g \rangle$ with M_i the set of messages of player i, and $g : \times_{i \in N} M_i \to X$ the allocation rule. Let $M := \times_{j \in N} M_j$ and $M_{-i} := \times_{j \in N \setminus \{i\}} M_j$, with m and m_{-i} generic elements, respectively.

A mechanism $\langle (M_i)_{i \in N}, g \rangle$ together with a state θ induce a strategic-form game $G(\theta)$ as follows. There is a set N of n players. The set of pure actions of player i is M_i, and player i's payoff when he plays m_i and his opponents play m_{-i} is $u_i(g(m_i, m_{-i}), \theta)$. Let Σ_i be the set of mixed strategies of player i and denote $\Sigma_{-i} := \times_{j \in N \setminus \{i\}} \Sigma_j$, with generic elements σ_i and σ_{-i}, respectively.

The aim of this paper is to study the problem of full implementation when players behave according to the solution concept of ε-minimax regret equilibrium (Renou and Schlag (2009)). Before defining our solution concept, we wish to stress that the concern for minimizing maximal regret does not arise from any behavioral or emotional considerations. Rather, it is a consequence of relaxing some of the axioms of subjective expected utility (e.g., the axiom of independence to irrelevant alternatives); we refer the reader to Hayashi (2008) or Stoye (2008) for recent axiomatizations. At state θ, player i's ex-post regret $r_i((m_i, m_{-i}), \theta)$ associated with the profile of messages (m_i, m_{-i}) is given by

$$r_i((m_i, m_{-i}), \theta) = \sup_{m_i' \in M_i} u_i(g(m_i', m_{-i}), \theta) - u_i(g(m_i, m_{-i}), \theta). \tag{1}$$

Player i's regret is thus the difference between player i's payoff when the profile of messages (m_i, m_{-i}) is played and the highest payoff he would have got, had he known that his opponents were playing m_{-i}. With a slight abuse of notation, we write $r_i((\sigma_i, \sigma_{-i}), \theta)$ for the expected regret when the profile (σ_i, σ_{-i}) of mixed strategies is played. A profile of strategies $(\sigma_i^*, \sigma_{-i}^*)$ is an ε-minimax regret equilibrium of the game $G(\theta)$ if

$$(1 - \varepsilon) r_i((\sigma_i^*, \sigma_{-i}^*), \theta) + \varepsilon \sup_{\sigma_{-i} \in \Sigma_{-i}} r_i((\sigma_i^*, \sigma_{-i}), \theta) \leq$$

$$(1 - \varepsilon) r_i((\sigma_i', \sigma_{-i}^*), \theta) + \varepsilon \sup_{\sigma_{-i} \in \Sigma_{-i}} r_i((\sigma_i', \sigma_{-i}), \theta)$$

for all $\sigma_i' \in \Sigma_i$, for all $i \in N$.

Let us briefly comment on our solution concept. In an ε-minimax regret equilibrium $(\sigma_i^*, \sigma_{-i}^*)$, player i believes that his opponents are playing σ_{-i}^* with probability $1 - \varepsilon$ and is completely uncertain about the play of his opponents, otherwise. And whenever uncertain about the play of his opponents, player i conjectures that his opponents play so as to maximize his regret. In particular, this implies that there is no mutual belief in rationality (unless $\varepsilon = 0$). To see this, suppose that there are only two players, 1 and 2, that player 2 has a strictly dominant action and player 1 believes with probability one that player 2 is rational. Since player 1 is certain that player 2 is rational, player 1 must conjecture that player 2 plays his strictly dominant action with probability one. However, the strictly dominant action of player 2 might not coincide with the action that maximizes player 1's regret. According to our solution concept, player 1 might therefore conjecture that player 2 plays a strictly dominated action with strictly positive probability (at most ε, however), a contradiction with his belief about player 2's rationality. (See Example 1.) The parameter ε quantifies the beliefs in rationality: the higher ε, the more doubtful a player is about the rationality of his opponents. Similarly, there is no common belief in conjectures (unless $\varepsilon = 0$). Two additional remarks are in order. Firstly, the concept of ε-minimax regret equilibrium with $\varepsilon = 0$ coincides with the concept of Nash equilibrium. Secondly, if $\varepsilon = 1$, there is no strategic considerations: a player simply minimizes his maximal regret, without making inferences about the play of his opponents, given his knowledge of their payoffs and the fact that they minimize maximal regret. We refer the reader to Renou and Schlag (2009) for additional results.

Definition 1 *Fix $\varepsilon \in [0, 1]$. The mechanism $\langle (M_i)_{i \in N}, g \rangle$ implements the social choice correspondence f in ε-minimax regret equilibrium if for all $\theta \in \Theta$, the following two conditions hold:*

(i) *For each $x \in f(\theta)$, there exists an ε-minimax regret equilibrium m^* of $G(\theta)$ such that $g(m^*) = x$.*

(ii) *If σ is an ε-minimax regret equilibrium of $G(\theta)$, then $g(m) \in f(\theta)$ for all m in the support of σ.*

Before proceeding, it is important to note that our definition of implementation in ε-minimax regret equilibrium follows Maskin (1999)'s definition and, in particular, considers mixed strategies in part (ii). This contrasts with most of the literature on

Nash implementation: with the notable exceptions of Maskin (1999) and Mezzetti and Renou (2009), mixed strategies are not considered in this literature.[2] With the concept of ε-minimax regret equilibrium, however, it is fundamental to consider mixed strategies: unlike the concept of Nash equilibrium, a player might be indifferent between two pure actions and yet strictly prefers a mixture of the two over each one.[3] We now illustrate our concept of implementation in ε-minimax regret equilibrium with the help of a simple example.

Example 1. There are two players, 1 and 2, two states of the world, θ and θ', and four alternatives, a, b, c, and d. The utilities are represented in the table below. For instance, at state θ, player 1's payoff of a is 10, while player 2's payoff is 2.

	θ		θ'	
	1	2	1	2
a	10	2	5	5
b	4	5	1	0
c	5	0	0	1
d	5	1	0	0

Consider the social choice function f with $f(\theta) = d$ and $f(\theta') = a$. It is implementable in Nash equilibrium (Maskin (1999)), as well as in rationalizable outcomes (Bergemann, Morris and Tercieux (2009)). A simple mechanism to implement f is the following: each player i has two messages m_i and m_i', and the allocation rule is given in the table below.

	m_2	m_2'
m_1	a	b
m_1'	c	d

In particular, at state θ, the strategic-form game $G(\theta)$ is:

---

[2]Note a substantial difference between Maskin (1999) and Mezzetti and Renou (2009): in part (i) of the definition of implementation, Maskin does not consider mixed Nash equilibria, while Mezzetti and Renou do.

[3]This follows from the axiomatization of the minimax regret criterion, which requires the uncertainty axiom. The uncertainty axiom says that if an individual is indifferent between two acts, he prefers a mixture of these acts over each of them. (See Stoye (2008).)

	m_2	m_2'
m_1	$10, 2$	$4, 5$
m_1'	$5, 0$	$5, 1$

So, if player 1 is certain that player 2 is rational, then (m_1', m_2') is the unique rationalizable outcome, hence the unique Nash equilibrium. Suppose now that player 1 has some doubts about the rationality of player 2 and, thus, is not certain that player 2 plays his strictly dominant action m_2'. Assume that player 1 believes that player 2 is rational with probability $1 - \varepsilon$ and, thus, believes that player 2 plays m_2' with probability at least $1 - \varepsilon$. With probability ε, player 1 is completely uncertain about the play of player 2 and conjectures that player 2 plays so as to maximize his (player 1) regret. We argue that (m_1', m_2') is nonetheless the unique ε-minimax regret equilibrium whenever ε is small enough.

Clearly, it is optimal for player 2 to play his strictly dominant action m_2'. Now, if player 1 considers playing m_1', his maximal regret is 5ε. With probability $1 - \varepsilon$, player 1 believes that player 2 is rational and thus plays m_2', in which case m_1' is player 1's best reply. With probability ε, player 1 conjectures that player 2 maximizes his (player 1) regret, and thus plays m_2. In that case, m_1 (and not m_1') is the best reply, so that the regret is 5.[4] Alternatively, if player 1 considers playing m_1, his maximal regret is 1. We might then conclude that m_1' minimizes player 1's maximal regret whenever $\varepsilon \leq 1/5$. However, if $\varepsilon > 1/6$, player 1 can guarantee himself a maximal regret of $5/6$ by randomizing between m_1 and m_1' with probability $5/6$ and $1/6$, respectively. In other words, m_1' is not a profitable deviation whenever $1/5 > \varepsilon > 1/6$, but the randomization between m_1 and m_1' is; this highlights the importance of considering mixed strategies. In fact, (m_1', m_2') is the unique ε-minimax regret equilibrium whenever $\varepsilon \leq 1/6$.[5] This simple example suggests that social choice correspondences that are Maskin monotonic and satisfies the no-veto power condition are implementable in ε-minimax regret equilibrium whenever ε is small enough, i.e., whenever the doubt about the rationality of others is small enough. Surprisingly, this turns out to be true even for arbitrarily large ε, i.e., for all $\varepsilon < 1$.

---

[4]That is $u_1(g(m_1, m_2), \theta) - u_1(g(m_1', m_2), \theta) = 5$.

[5]If $\varepsilon > 1/6$, there is a unique ε-minimax regret equilibrium, in which player 1 plays m_1 with probability $5/6$ and player 2 plays m_2' with probability one.

3 Monotonicity and no-veto power

The main result of this note is that social choice correspondences that are Maskin monotonic and satisfy the no-veto power condition are implementable in ε-minimax regret equilibrium for all $\varepsilon \in [0, 1)$. Maskin's result is surprisingly robust: even arbitrarily large uncertainties about the rationality of opponents do not undermine the implementation of social choice correspondences in $\mathcal{F}_{NE}$.

Proposition 1 *Let $n \geq 3$. If the social choice correspondence f is Maskin monotonic and satisfies no-veto power, then it is implementable in ε-minimax regret for all $\varepsilon \in [0, 1)$.*

The intuition for Proposition 1 is simple. The canonical mechanism for Nash implementation features *an integer game*, whereby whenever it is not the case that at least $n-1$ players announce the same message, the alternative implemented is the alternative nominated by the player announcing the highest integer (if there are several such players, choose the player with the lowest index). Consider now a profile of messages such that the integer game applies. The regret to any player is then the difference in payoffs between his most preferred alternative and the alternative implemented. Indeed, had the player known the integers of the others, he could have chosen a strictly higher integer and got his most preferred alternative. Clearly, his maximal regret is then the difference in payoffs between his most preferred alternative and his less preferred alternative. Furthermore, regardless of the strategy a player follows, his opponents can always trigger the integer game and impose to the player his maximal regret. With probability ε, the maximal regret to a player is therefore constant and, consequently, only the regret of facing rational opponents matters. In turn, this implies Nash behavior.

Proof Since f is Maskin monotonic and satisfies the no-veto power condition, it is implementable in Nash equilibrium. In particular, the mechanism of Maskin and Sjöström (2002) implements f. The mechanism has the following distinctive feature: each player i has to report an alternative x^i, a state θ^i, an integer z^i and a mapping $\alpha^i : X \times \Theta \to X$ from alternatives and states to alternatives satisfying $\alpha^i(x, \theta) \in L_i(x, \theta)$ for all (x, θ). Moreover, under rule 3 of the mechanism, whenever it is not the case that at least $n-1$ players announce the same alternative x, the same state θ and the integer 1, the alternative implemented is the alternative announced by the player with the highest integer (if there are several such players, choose the player with the lowest index): it is

an integer game.

We claim that the maximal regret a player can experience is independent of the strategy σ_i he plays. To see this, assume that the state of the world is θ. Suppose that the profile (m_i, m_{-i}) is realized such that rule 3 applies, i.e., $g((m_i, m_{-i})) = x^{i^*}$ with i^* the player with the lowest index among the players announcing the highest integer. The regret to player i in state θ is therefore:

$$\max_{x \in X} u_i(x, \theta) - u_i(g(m_i, m_{-i}), \theta),$$

since player i can always get his most preferred alternative by announcing an integer strictly higher than the ones of his opponents. Furthermore, by choosing m_{-i} such that at least one player other than i announces an integer higher that the one of player i and all players other than i announce the less preferred alternative of player i (i.e., $m_{-i} = (\theta^j, \alpha^j, x^j, z^j)_{j \in N \setminus \{i\}}$ with $x^j \in \min_{x \in X} u_i(x, \theta)$ for all $j \in N \setminus \{i\}$ and $z^j > z^i$ for at least two players j), the maximal regret to player i is

$$\max_{x \in X} u_i(x, \theta) - \min_{x \in X} u_i(x, \theta) := K_i(\theta).$$

It then follows that for any σ_i, the supremum of the regret is attained in $K_i(\theta)$. (It is actually a maximum whenever σ_i has a finite support. But no maximum exists when σ_i has an unbounded support.)

Finally, let $(\sigma_i^*, \sigma_{-i}^*)$ be a Nash equilibrium of the mechanism at state θ. Note that this is equivalent to $r_i((\sigma_i^*, \sigma_{-i}^*), \theta) \leq r_i((\sigma_i, \sigma_{-i}^*), \theta)$ for all $\sigma_i \in \Sigma_i$. It follows that

$$(1 - \varepsilon) r_i((\sigma_i^*, \sigma_{-i}^*), \theta) + \varepsilon \sup_{\sigma_{-i} \in \Sigma_{-i}} r_i((\sigma_i^*, \sigma_{-i}), \theta) =$$
$$(1 - \varepsilon) r_i((\sigma_i^*, \sigma_{-i}^*), \theta) + \varepsilon K_i(\theta) \leq$$
$$(1 - \varepsilon) r_i((\sigma_i, \sigma_{-i}^*), \theta) + \varepsilon K_i(\theta) =$$
$$(1 - \varepsilon) r_i((\sigma_i, \sigma_{-i}^*), \theta) + \varepsilon \sup_{\sigma_{-i} \in \Sigma_{-i}} r_i((\sigma_i, \sigma_{-i}), \theta),$$

for all $\sigma_i \in \Sigma_i$, so that $(\sigma_i^*, \sigma_{-i}^*)$ is an ε-minimax regret equilibrium.

Conversely, if $(\sigma_i^*, \sigma_{-i}^*)$ is not a Nash equilibrium, there exist a player i and a strategy σ_i' such that $u_i((\sigma_i', \sigma_{-i}^*), \theta) > u_i((\sigma_i^*, \sigma_{-i}^*), \theta)$, so that $r_i((\sigma_i^*, \sigma_{-i}^*), \theta) > r_i((\sigma_i', \sigma_{-i}^*), \theta)$. It follows from the above arguments that $(\sigma_i^*, \sigma_{-i}^*)$ is not an ε-minimax regret equilibrium provided that $\varepsilon < 1$. This completes the proof. $\qquad\square$

Conversely, if a social choice correspondence is implementable in ε-minimax regret equilibrium for all $\varepsilon \in [0, 1)$, then it is Maskin monotonic. Maskin monotonicity is thus necessary and almost sufficient for implementation in ε-minimax regret equilibrium for all $\varepsilon \in [0, 1)$.

4   Discussion

We conclude this note with two remarks. Firstly, we show with the help of a simple example that Maskin monotonicity is not a necessary condition for implementation in ε-minimax regret equilibrium, even when implementation is required for all $\varepsilon \in (0, 1)$.[6]

Example 2. Consider an economy with a perfectly divisible good in fixed supply of one unit. There are three consumers in this economy. Consumers 1 and 2 have strictly monotonic preferences, while consumer 3 has single-peaked preferences. In state θ, consumer 3's preferences are single-peaked at $1/2$, while they are single-peaked at $1/3$ in state θ'. In both states, consumer 3 strictly prefers a positive consumption over a zero consumption. The social planner aims to implement the allocation $f(\theta) = \{(1/3, 1/3, 1/3)\}$ at state θ and the allocation $f(\theta') = \{(1/4, 1/4, 1/2)\}$ at state θ'.

The social choice function f cannot be implemented in Nash equilibrium. Intuitively, consumer 3 prefers the allocation $f(\theta')$ in state θ and the allocation $f(\theta)$ in state θ', so that he has no incentive to truthfully reveal his preferences. Indeed, the social choice function f is not Maskin monotonic, a necessary condition for Nash implementation.

Yet, the social choice function f is implementable in ε-minimax regret equilibrium for all $\varepsilon \in (0, 1)$: there is a discontinuity at $\varepsilon = 0$. To see this, consider the following mechanism: Each consumer has to announce either a state θ or θ', or an integer z. If all consumers announce a state, the allocation rule is as below:

	θ	θ'			θ	θ'
θ	$1/3, 1/3, 1/3$	$1/3, 0, 1/2$		θ	$0, 0, 1/3$	$0, 1/3, 1/3$
θ'	$0, 1/3, 1/2$	$0, 0, 1/2$		θ'	$1/3, 0, 1/3$	$1/4, 1/4, 1/2$

$$\theta \qquad\qquad\qquad\qquad\qquad \theta'$$

Alternatively, if exactly one consumer announces an integer and the two others announce a state, the allocation is the bundle $(0, 1/3, 0)$ if consumer 1 announces the integer,

[6]The example is adapted from Sjötröm (1994).

$(1/3, 0, 0)$ if consumer 2 announces the integer and $(0, 0, 0)$ if consumer 3 announces the integer. If two consumers or more announce an integer, the allocation is the bundle $(1/3, 0, 0)$ if consumer 1 announces a higher integer than consumer 2 or consumer 2 announces a state, and $(0, 1/3, 0)$ if consumer 2 announces a strictly higher integer than consumer 1 or consumer 1 announces a state.

We have to show that at state θ (resp., θ'), the profile of messages (θ, θ, θ) (resp., $(\theta', \theta', \theta')$) is the unique ε-minimax regret equilibrium for any $\varepsilon \in (0, 1)$. Suppose that the true state is θ. We first argue that in any ε-minimax regret equilibrium at state θ, consumer 3 announces the true state θ. To see this, observe that the message θ is weakly dominant at state θ for consumer 3 and, consequently, consumer 3's regret of playing θ is strictly smaller (in fact, zero) than the regret of playing any other strategies, regardless of the strategies of consumers 1 and 2. We next argue that consumers 1 and 2's maximal regret is minimized at θ when consumer 3 announces the true state θ (with probability $1 - \varepsilon$). To see this, note that, conditional on consumer 3 announcing the true state, it is weakly dominant for consumers 1 and 2 to match the announcement of consumer 3. Moreover, the presence of an integer game between consumers 1 and 2 implies, as with Proposition 1, that the maximal regret to consumer 1 (resp., consumer 2) when facing "irrational" opponents is independent of the strategy he plays. Consequently, (θ, θ, θ) is the unique ε-minimax regret equilibrium at state θ. A similar reasoning applies at state θ'.[7]

In addition, it is worth noting that the social choice function f is implementable in undominated Nash equilibrium (Palfrey and Srivastava (1991), Jackson et al. (1994) and Sjöström (1994)). This suggests a possible connection between implementation in ε-minimax regret equilibrium and in undominated Nash equilibrium (Palfrey and Srivastava (1991)). In particular, we conjecture that if a social choice correspondence is implementable in ε-minimax regret equilibrium, then it is implementable in undominated Nash equilibrium. Unfortunately, we have not been able to prove it. Nonetheless, the following example due to Jackson (1992) shows that both concepts do not coincide.

Example 3. There are five players, labeled 1 to 5, two states of the world θ and θ'

---

[7]The logic of this example can be generalized to show that, in separable environments (Jackson et al. (1994)), any individually rational social choice function is implementable in ε-minimax regret equilibrium for all $\varepsilon \in (0, 1)$.

and two alternatives x and y. The preferences are:

	θ					θ'			
1	2	3	4	5	1	2	3	4	5
y	y	x	x	y	y	y	x	x	x
x	x	y	y	x	x	x	y	y	y

The social choice correspondence selects the less preferred alternative of player 5 at each state, i.e., $f(\theta) = \{x\}$ and $f(\theta') = \{y\}$. It is implementable in undominated Nash equilibrium, but not in ε-minimax regret equilibrium for any ε. Assume by contraction that f is implementable in ε-minimax regret equilibrium by the mechanism $\langle M, g \rangle$. At state θ, there must exist a *pure ε-minimax regret equilibrium* m^* with $g(m^*) = x$. We argue that for m^* to be a pure equilibrium, it must be that $u_i(g(m_i^*, m_{-i}), \theta) \geq u_i(g(m_i, m_{-i}), \theta)$ for all m_i, for all m_{-i}, for all $i \in \{1, 2, 5\}$. Consider player 5. (A similar argument applies to players 1 and 2.) To the contrary, assume that there exists a profile of messages (m_5', m_{-5}') such that $g(m_5^*, m_{-5}') = x$ and $g(m_5', m_{-5}') = y$. The maximal regret of playing m_5^* is therefore at least $\varepsilon(u_5(y, \theta) - u_5(x, \theta))$. It is $\varepsilon(u_5(y, \theta) - u_5(x, \theta))$ whenever $g(m_5, m_{-5}^*) = x$ for all $m_5 \in M_5$, and $u_5(y, \theta) - u_5(x, \theta)$, otherwise. Clearly, if m_5' weakly dominates m_5^*, m^* cannot be an equilibrium. Similarly, if $m_{-5}' = m_{-5}^*$, m^* cannot be an equilibrium: the regret of playing m_5^* would be $u_5(y, \theta) - u_5(x, \theta)$, while the regret of playing m_5' would be at most $\varepsilon(u_5(y, \theta) - u_5(x, \theta))$. However, if none of the previous holds, it has to be that the maximal regret of playing m_5' is the same of playing m_5^*, i.e., $\varepsilon(u_5(y, \theta) - u_5(x, \theta))$. But then player 5 has a profitable deviation: he can randomize between m_5^* and m_5' and strictly decreases his maximal regret. A direct implication of the above result is that $g(m_1, m_2, m_3^*, m_4^*, m_5) = x$ for all m_1, for all m_2, for all m_5. Similarly, at state θ', there must exist a *pure ε-minimax regret* equilibrium m^{**} with $g(m^{**}) = y$. Following the same logic as above, it follows that $g(m_1^{**}, m_2^{**}, m_3, m_4, m_5) = y$ for all m_3, for all m_4, for all m_5, a contradiction.

References

[1] Dirk Bergemann, Stephen Morris and Olivier Tercieux, 2009, "Rationalizable Implementation," Mimeo, Yale University.

[2] Antonio Cabrales and Roberto Serrano, 2008, Implementation in Adaptive Better-Response Dynamics, Mimeo, Brown University.

[3] Matthew O. Jackson, 1992, "Implementation in Undominated Nash Strategies: a Look at Bounded Mechanisms," Review of Economic Studies, 59, pp. 461-478.

[4] Matthew O. Jackson, 2001, "A Crash Course in Implementation Theory," Social Choice and Welfare, 18, pp. 655-708.

[5] Matthew O. Jackson, Thomas R. Palfrey and Sanjay Srivastava, 1994, "Undominated Nash Implementation in Bounded Mechanisms," Games and Economic Behavior, 6, pp. 474-501.

[6] Takashi Hayashi, 2008, "Regret aversion and opportunity dependence," Journal of Economic Theory, 139, pp. 242-268.

[7] Eric Maskin, 1999, "Nash Equilibrium and Welfare Optimality," Review of Economic Studies, 66, pp. 23-38.

[8] Eric Maskin and Tomas Sjöström, 2002, "Implementation Theory," in Amartya Sen and Kotaro Suzumura (eds), Handbook of Social Choice and Welfare, vol. 1, Amsterdam: North Holland, pp. 237-288.

[9] Claudio Mezzetti and Ludovic Renou, 2009, "Implementation in Mixed Nash Equilibrium," Working Paper 09/10, University of Leicester.

[10] Thomas R. Palfrey and Sanjay Srivastava, 1991, "Nash implementation using Undominated Strategies," Econometrica, 59, pp. 479-501

[11] Ludovic Renou and Karl H. Schlag, 2009, "Minimax Regret and Strategic Uncertainty," forthcoming in Journal of Economic Theory.

[12] Tomas Sjöström, 1994, "Implementation in Undominated Nash Equilibria without Integer Games," Games and Economic Behavior, 6, pp. 502–511

[13] Joerg Stoye, "Statistical Decisions under Ambiguity," Mimeo, New York University, 2008.

[14] Norovsambuu Tumennasan, "To Err is Human: Implementation in Quantal Response Equilibria," Mimeo, Brown University, 2008