



Computer Science and Artificial Intelligence Laboratory
Technical Report

MIT-CSAIL-TR-2013-029

December 3, 2013

Bridging Utility Maximization and Regret Minimization
Alessandro Chiesa, Silvio Micali, and Zeyuan
Allen Zhu

Bridging Utility Maximization and Regret Minimization

Alessandro Chiesa

alexch@csail.mit.edu

MIT CSAIL

Silvio Micali

silvio@csail.mit.edu

MIT CSAIL

Zeyuan Allen Zhu

zeyuan@csail.mit.edu

MIT CSAIL

December 3, 2013

1 Introduction

Rational players have been modeled in two main ways.

- A utility-maximizing player \mathcal{U} eliminates all his dominated strategies to compute his set of undominated ones, UD . Notice that \mathcal{U} cannot further refine UD based on utility maximization. If UD consists of a single strategy s (necessarily a dominant one), then \mathcal{U} of course chooses s . *But, if UD contains multiple strategies, which one should \mathcal{U} choose?*
- A regret-minimizing player \mathcal{R} eliminates all his non regret-minimizing strategies so as to compute his set of regret-minimizing strategies, RM . He might even continue this process k times, until he is satisfied or no further elimination is possible. Let us denote the final set of strategies he obtains this way by RM^k . If RM^k consists of a single strategy s , he of course chooses s . *But, if RM^k contains multiple strategies, which one should \mathcal{R} choose?*

In both cases, “a random strategy” or “the lexicographic first strategy” are certainly possible answers. But another answer is that, when he is ‘no longer able to apply his favorite way of reasoning’, even a die-hard utility maximizer \mathcal{U} will resort to regret minimization to refine UD , and even a die-hard regret minimizer \mathcal{R} will resort to utility maximization to refine RM^k . In principle, the two final sets of strategies obtained by such different refinement procedures could be vastly different. Our next structural theorem, however, guarantees that they coincide.

Abusing notation a bit, consider UD and RM also to be “operators” acting on sets of strategies. In this case $\text{UD}(\text{UD}) = \text{UD}$, while $\text{RM}^2 \stackrel{\text{def}}{=} \text{RM}(\text{RM})$ may be a strict subset of RM . Then, we prove that, whether or not the players are Knightian,

Theorem 1 (Informal). *The set of strategies obtained after applying, in arbitrary order, k times the operator RM and at least once the operator UD coincides with $\text{RM}^k \cap \text{UD}$.*

For instance, $\text{RM}(\text{RM}(\text{UD}(\text{RM}(\text{RM}(\text{UD})))) = \text{RM}^4(\text{UD}) = \text{RM}^4 \cap \text{UD}$.

Whether players are utility maximizers or regret minimizers is an old question. In particular, iterated regret minimization (using beliefs) has recently been advocated as a valid solution concept, indeed the only one capable of explaining the actual behavior of the players in some settings [HP12].

Theorem 1 has an immediate but significant consequence for mechanism design. Namely,

*For all mechanisms M and social choice correspondences f ,
if M implements f in RM strategies or in UD strategies,
then M is automatically guaranteed to implement f also in RM(UD) strategies.¹*

2 Preliminaries

We prove our theorem in the language of decision theory: namely, for a single player “against Nature”. Results for n -player (strategic or pre-Bayesian) games follow as corollaries. This is because the definitions of dominance and regret are universally quantified over other players’ strategies, (and the player’s true valuation in his candidate set if it is the Knightian case,) which can be treated as Nature’s strategies.

Let \mathcal{S} be a compact set of strategies of a player, and T a compact set of states of Nature.² We denote by U the (continuous) utility function of the player, where $U(s, t)$ is the utility under strategy $s \in \mathcal{S}$ when Nature’s state is $t \in T$. Regret-minimizing strategies and undominated strategies are defined as follows:

- Given a *menu* $S \subseteq \mathcal{S}$ of strategies, the player’s (maximum) regret for a strategy $s \in S$ in menu S , denoted by $R_S(s)$, is the maximum difference, taken over all

¹Indeed, for $i = 1$ Theorem 1 implies that $\text{RM(UD)} = \text{RM} \cap \text{UD} \subseteq \text{RM}$. Of course, to enforce the same guarantee one could just demand that M implements f in $\text{RM} \cup \text{UD}$ strategies, but this is a very strong demand. Indeed $\text{RM} \cup \text{UD}$ could be a much larger set than $\text{RM} \cap \text{UD}$.

²For instance, in the Knightian setting of the VCG (see Section ??), when analyzing a player i , \mathcal{S} consists of all possible bidding strategies of player i , and T is the cartesian product of (1) all possible bidding strategy sub-profiles of i ’s opponents and (2) all possible true valuations of player i in his set K_i .

Both \mathcal{S} and T may be infinite, and \mathcal{S} may be convex in order to allow arbitrary mixed strategies to be considered.

possible Nature's states $t \in T$, between the utility the player gets by playing s , and that he could have gotten by "best responding" to t ; formally, $R_S(s) \stackrel{\text{def}}{=} \max_{t \in T} \left(\max_{s^* \in S} U(s^*, t) - U(s, t) \right)$.

Therefore, the regret-minimizing strategies with respect to a menu $S \subseteq \mathcal{S}$, denoted by $\text{RM}(S)$, is the set of strategies that minimize the regret: $\text{RM}(S) \stackrel{\text{def}}{=} \arg \min_{s \in S} R_S(s)$.

- Given two strategies $s, s' \in \mathcal{S}$, by definition s' weakly dominates s , denoted by $s' \succ s$, if

$$\forall t \in T, U(s', t) \geq U(s, t) \quad \text{and} \quad \exists t \in T, U(s', t) > U(s, t) .$$

Given a menu $S \subseteq \mathcal{S}$ of strategies, the player's undominated strategies consist of those that are not weakly dominated by any weakly undominated strategy.³ Formally,

$$\begin{aligned} \text{UD}(S) &\stackrel{\text{def}}{=} S \setminus \{s \in S : \exists s' \in S \text{ s.t. } (s' \succ s) \wedge (\nexists s'' \in S, s'' \succ s')\} \\ &= \{s \in S : \nexists s' \in S \text{ s.t. } (s' \succ s) \wedge (\nexists s'' \in S, s'' \succ s')\} \end{aligned}$$

We now state two simple facts which follow easily from the above definitions:

Fact 2.1. *For any menu $\tilde{S} \subseteq \mathcal{S}$,*

- (a) *if $s \prec s'$ for some $s, s' \in \tilde{S}$, then $R_{\tilde{S}}(s) \geq R_{\tilde{S}}(s')$, and*
- (b) *the regret values of a strategy with respect to \tilde{S} and $\text{UD}(\tilde{S})$ are the same,*

³In many cases of interest (e.g., when the set of pure strategies is finite, or when the mechanism is the VCG), weakly undominated strategies coincide with undominated ones, and this is why we directly adopted that simpler notion in Section ?? for Knightian auctions. As argued by Jackson [Jac92], however, the above level of precision is required when handling the general case. In particular, it may happen that every pure strategy is weakly dominated by another one in an infinite chain, and in such a case all strategies are undominated but weakly dominated.

namely:⁴

$$R_{\tilde{S}}(s) = \max_{t \in T} \left(\max_{s^* \in \tilde{S}} U(s^*, t) - U(s, t) \right) = \max_{t \in T} \left(\max_{s^* \in \text{UD}(\tilde{S})} U(s^*, t) - U(s, t) \right) = R_{\text{UD}(\tilde{S})}(s) .$$

3 Formal Statement and Proof of Our Theorem

Established our language, we prove our theorem as a corollary of the following lemma.

Lemma 1. *For any menu $S \subseteq \mathcal{S}$, $\text{UD}(\text{RM}(S)) = \text{RM}(\text{UD}(S)) = \text{RM}(S) \cap \text{UD}(S)$.*

Proof. We divide the proof into six steps:

1. $\text{RM}(\text{UD}(S)) \subseteq \text{RM}(S)$.

For any $s \in \text{RM}(\text{UD}(S))$, we show that $s \in \text{RM}(S)$ by proving that s has minimum regret among all strategies in S . Indeed:

- For any other strategy $s' \in \text{UD}(S)$, it holds that $R_{\text{UD}(S)}(s) \leq R_{\text{UD}(S)}(s')$. By Fact 2.1b, we deduce that $R_S(s) \leq R_S(s')$.
- For any other strategy $s' \in S \setminus \text{UD}(S)$, it holds that $s' \prec s''$ for some $s'' \in \text{UD}(S)$ and $R_S(s) \leq R_S(s'')$. By Fact 2.1a, we deduce that $R_S(s) \leq R_S(s'') \leq R_S(s')$.

2. $\text{RM}(\text{UD}(S)) \subseteq \text{UD}(\text{RM}(S))$.

Given that $\text{RM}(\text{UD}(S)) \subseteq \text{RM}(S)$ (proved above), if there is some $s \in \text{RM}(\text{UD}(S))$ with $s \notin \text{UD}(\text{RM}(S))$, then s must be weakly dominated by some other strategy $s' \in \text{RM}(S)$, namely $s \prec s'$, but s' cannot be weakly dominated by any other strategy in $\text{RM}(S)$, by definition of UD.

Now we show that s' cannot be weakly dominated by any strategy in S as well.

Suppose not, that is $s' \prec s''$ where $s'' \in S$. Then $s'' \notin \text{RM}(S)$ as we have just

⁴The equality in the middle is since any strategy $s^* \in \tilde{S} \setminus \text{UD}(\tilde{S})$ must be weakly dominated by some $s^{**} \in \tilde{S}$, giving at least as good utilities as s^* for *any* $t \in T$. Therefore, such choices of s^{**} can be ignored in the inner max.

argued. However, using Fact 2.1a we have $R_S(s') \geq R_S(s'')$, implying that $s'' \in \text{RM}(S)$ since $s' \in \text{RM}(S)$, giving a contradiction to $s'' \notin \text{RM}(S)$.

In sum, we showed that s is weakly dominated by $s' \in S$, and in addition s' cannot be weakly dominated by any strategy in S , contradicting the fact that $s \in \text{UD}(S)$.

3. $\text{UD}(\text{RM}(S)) \subseteq \text{UD}(S)$.

Suppose not, that is, there exists some $s \in \text{UD}(\text{RM}(S))$ that is not in $\text{UD}(S)$. By the definition of $\text{UD}(S)$, the strategy s must be weakly dominated by some $s' \in S$, and in addition s' cannot be weakly dominated by any other strategy in S . There are two cases here.

- The first case is when $s' \in \text{RM}(S)$. This case is impossible because $s \in \text{UD}(\text{RM}(S))$ implies that if s is weakly dominated by $s' \in \text{RM}(S)$, then s' must also be weakly dominated, contradicting the fact that s' cannot be weakly dominated by any strategy in S .
- The second case is when $s' \notin \text{RM}(S)$. Since $s \prec s'$, by Fact 2.1a we have $R_S(s) \geq R_S(s')$. However, because $s \in \text{UD}(\text{RM}(S))$ implies that $s \in \text{RM}(S)$, it must hold that s' is a regret minimizer with respect to S , contradicting the fact that $s' \notin \text{RM}(S)$.

4. $\text{UD}(\text{RM}(S)) \subseteq \text{RM}(\text{UD}(S))$.

Given that $\text{UD}(\text{RM}(S)) \subseteq \text{UD}(S)$ (proved above), consider any strategy $s \in \text{UD}(\text{RM}(S))$, and suppose that $s \notin \text{RM}(\text{UD}(S))$. Then there exists some $s' \in \text{UD}(S)$ satisfying $R_{\text{UD}(S)}(s) > R_{\text{UD}(S)}(s')$. This implies, through Fact 2.1b, that $R_S(s) > R_S(s')$, contradicting the fact that $s \in \text{RM}(S)$.

5. $\text{RM}(\text{UD}(S)) \subseteq \text{RM}(S) \cap \text{UD}(S)$.

Trivial given the previous steps: $\text{RM}(\text{UD}(S)) \subseteq \text{UD}(S)$ and $\text{RM}(\text{UD}(S)) = \text{UD}(\text{RM}(S)) \subseteq \text{RM}(S)$.

6. $\text{RM}(S) \cap \text{UD}(S) \subseteq \text{RM}(\text{UD}(S))$.

Take any strategy $s \in \text{RM}(S) \cap \text{UD}(S)$, and suppose that $s \notin \text{RM}(\text{UD}(S))$. Then there exists some $s' \in \text{UD}(S)$ satisfying $R_{\text{UD}(S)}(s) > R_{\text{UD}(S)}(s')$. This implies, through Fact 2.1b, that $R_S(s) > R_S(s')$, contradicting the fact that $s \in \text{RM}(S)$. \square

It is not hard to see that Lemma 1 implies our theorem. That is,

Theorem 1 (restated). *From any menu $S \subseteq \mathcal{S}$, a player who applies, in arbitrary order, i times the operator RM and at least once the operator UD , always obtains the same set of surviving strategies:*

$$\text{RM}^i(S) \cap \text{UD}(S) .$$

4 Pure vs. Mixed Strategies

So far we have been ambiguous, when discussing undominated strategies and regret-minimizing ones, about whether or not the players consider only pure strategies or also mixed ones. When only pure strategies are allowed, a utility maximizer compares only between his pure strategies for the notion of dominance and plays a pure undominated one, while a regret minimizer picks a pure strategy that minimizes regret among his pure strategies.

Our theorem and lemma are stated for pure strategies.

When mixed strategies are allowed, the definitions of UD and RM need more careful attention. It is easy to see that, when considering mixed strategies for regret minimizers, the only change needed is to allow such a minimizer to choose a mixed strategy that minimizes his expected regret among all his mixed ones (see e.g., [HB04, HP12]). Note that, it is easy to construct examples in which a mixed strategy yields strictly smaller regret than any pure strategy.

It is important to realize, however, that if we allow regret minimizers to consider mixed strategies, we *should* also allow utility maximizers to consider mixed strategies. For instance, our structural lemma (Lemma 1) would have difficulty to equate a set

of pure strategies and a set of mixed ones. A utility maximizer may consider mixed strategies when determining that a strategy s is weakly dominated by another strategy s' . The two interesting cases to consider are (1) s is pure and s' is mixed; and (2) both s and s' are mixed. Traditionally, most attention has been devoted to the first case, but the second has been studied too (see for instance [CS05, RS10]). Clearly, UD can be defined in both cases, and yields a more “refined” set of strategies in the second case.⁵ It is actually under this more refined case that our structural lemma holds. In a sense, we have nothing to lose and something to gain by adopting a more flexible definition, after all the right notions are those yielding the right theorems.

Comparison with the Notion of Hyafil and Boutilier. Hyafil and Boutilier [HB04] studied the notion of minimax-regret equilibrium in a setting where players have beliefs, and provided an LP-based solution for constructing mechanisms in certain restricted cases. At high level, they study a notion of regret that is based on beliefs about possible types of the opponents, and then consider a notion of equilibrium based on regret. Let us explain.

Hyafil and Boutilier assume that each player i forms a belief T_{-i} about his opponents’ possible types. Given this belief, player i can compute, for any strategy σ_i of his, strategy profile σ_{-i} of his opponents, and his own type θ_i , the maximum regret $R_i(\sigma_i, \sigma_{-i}, \theta_i)$ with respect to his opponents’ possible types in T_{-i} . Then, Hyafil and Boutilier define a minimax-regret Nash equilibrium to be a strategy profile $(\sigma_1, \dots, \sigma_n)$ in which no player can deviate to increase his maximum regret.

The above solution concept of minimax-regret equilibrium coincides with ours *when players have no belief about other players’ types*. Indeed, if a player has no prior knowledge about his opponents, the notion of maximum regret $R_i(\sigma_i, \sigma_{-i}, \theta_i)$ will not depend on σ_{-i} in general (at least for non-degenerate strategies whose range coincides

⁵Let UD^{pure} be the set of (pure) undominated strategies in the first case, and UD be the set of (possibly mixed) undominated strategies in the second case. Then, UD is a more “refined” notion of undominated strategies than UD^{pure} because $\text{UD}^{\text{pure}} \subseteq \text{UD} \subseteq \Delta(\text{UD}^{\text{pure}})$, i.e., UD^{pure} coincides with the support of UD. For this reason, there is no difference in choosing between the two notions in most of the literature (see [CS05, footnote 2]).

with all possible actions), and therefore in any minimax-regret Nash equilibrium a player simply chooses the strategy that minimizes maximum regret, like we do in this paper.

Comparison with the Notion of Renou and Schlag.

Renou and Schlag [RS10] also proposed a solution concept called minimax-regret equilibrium, with respect to possible beliefs about the other players' actions. They studied strategic games (in which each player only has a single type). Their solution concept does not coincide with Hyafil and Boutilier. In fact, although their solution concept is called an "equilibrium", the strategies of a player's opponents are considered all on the regret level, that is, they assume that a player always chooses a strategy minimizing the maximum regret over all possible strategies (according to his beliefs) of the his opponents. When players have no beliefs of their opponents, this notion trivially coincides with ours (after suitably generalizing it to allow players to have types).

References

- [CS05] Vincent Conitzer and Tuomas Sandholm. Complexity of (iterated) dominance. In *Proceedings of the 6th ACM conference on Electronic commerce, EC '05*, pages 88–97, New York, NY, USA, 2005. ACM.
- [HB04] Nathanael Hyafil and Craig Boutilier. Regret Minimizing Equilibria and Mechanisms for Games with Strict Type Uncertainty. In *Proceedings of the 20th conference on Uncertainty in artificial intelligence*, pages 268–277, July 2004.
- [HP12] Joseph Y. Halpern and Rafael Pass. Iterated regret minimization: A new solution concept. *Games and Economic Behavior*, 74(1):184–207, January 2012. A preliminary version appeared in IJCAI'09.
- [Jac92] Matthew O Jackson. Implementation in undominated strategies: A look at bounded mechanisms. *Review of Economic Studies*, 59(4):757–75, October 1992.
- [RS10] Ludovic Renou and Karl H. Schlag. Minimax regret and strategic uncertainty. *Journal of Economic Theory*, 145(1):264–286, January 2010.

