Regret Minimizing Equilibria of Games with Strict Type Uncertainty

Nathanael Hyafil and Craig Boutilier

Dept of Computer Science, U. of Toronto
{nhyafil,cebly}@cs.toronto.edu

Abstract

In the standard mechanism design setting, the type (e.g., utility function) of an agent is not known by other agents, nor is it known by the mechanism designer. When this uncertainty is quantified probabilistically, a mechanism induces a game of incomplete information among the agents. However, in many settings, uncertainty over utility functions cannot easily be quantified. We consider the problem of incomplete information games in which type uncertainty is *strict* or unquantified. We propose the use of *minimax regret* as a decision criterion in such games, a robust approach for dealing with type uncertainty. We define *minimax-regret equilibria* and prove that these exist in mixed strategies for finite games. We also briefly discuss mechanism design in this framework, with minimax regret as an optimization criterion for the designer itself, and the automated optimization of such mechanisms.

1 Minimax Regret

Minimax regret [10, 2] is a common criterion for decision making when uncertainty over consequences of decisions is not quantified probabilistically, a case we refer to as *strict uncertainty*.[1] Minimax regret is usually advocated when agents are assumed to behave in a non-Bayesian way, either by choice or because prior information is unavailable or too expensive to construct. In this context, we view minimax regret as more reasonable than alternatives such as the maximin (pessimistic), maximax (optimistic) or Hurwitz criteria, since it takes into account much more information than simply the worst and best cases.

More recently, minimax regret has been proposed as a suitable criterion for one agent (say, decision support software) to make decisions on behalf of another, when the first is uncertain about the true utility function of the second [13]. In such cases, regret provides a useful measure of the potential error or loss associated with a decision one makes on behalf of another (rather than relying on notions of lost opportunity) and can be used very effectively to drive the process of preference elicitation [13, 3]. Furthermore, it offers a very practical alternative to Bayesian approaches, since priors over the set of possible utility functions can be very difficult to obtain, represent, and update in response to queries and observations. In contrast, bounds or linear constraints on utility function parameters are much easier to obtain, update, and support computationally effective regret computation.

In this paper, we study the use of minimax regret in the context of games of incomplete information, where the (strict) uncertainty is about the types of the other agents. We will also discuss (briefly) how the use of minimax regret can be exploited in mechanism design.

[1] Regret has been used in other ways to provide alternatives to expected utility theory, even when distributional information is available (e.g., *regret theory* [6, 1]). This view could also be adopted in what follows, though we do not advocate this as a suitable normative decision criterion.

2 Games with Strict Type Uncertainty

An *incomplete information game with strict type uncertainty* consists of the same components as a (Bayesian) incomplete information game, but has a *qualitative* or *strict prior* rather than a probabilistic prior. Specifically, we assume an action set A_i, type space Θ_i, and utility function u_i for each agent i. We assume a *strict prior* $T \subseteq \Theta$ representing (common) beliefs about possible type profiles held by the agents in the game. Intuitively, Θ denotes the set of possible types from a "structural" perspective, while T denotes what is believed: only type profiles $\theta \in T$ are considered to be possible given the information possessed by the participants.[2]

As in the standard setting, we assume each agent knows its own type. Its beliefs about the types of other agents, given its type θ_i, is given by the set $T(\theta_i) = \{\theta_{-i} : \langle \theta_i, \theta_{-i} \rangle \in T\}$ (i.e., those type profiles consistent with its own known type). Strategies are defined in the usual way as mappings from types to (mixed) action choices: $\Theta_i \mapsto \Delta(A_i)$. Let Σ_i denote the set of (pure or mixed) strategies for agent i

2.1 Minimax-regret Equilibrium

Defining the *expected utility* of a fixed strategy σ_i requires some distribution over the possible types of other agents. Without distributional information, we must adopt some qualitative decision criterion to evaluate and compare strategies. Here we propose the use of the minimax-regret decision criterion.

Definition 1 The *regret* of strategy σ_i for agent i with type θ_i, given strategy profile σ_{-i} and type profile θ_{-i} of the other agents, is

$$R_i(\sigma_i | \theta_i, \theta_{-i}, \sigma_{-i}) = \max_{\sigma_i' \in \Sigma_i} u_i(\langle \sigma_i'(\theta_i), \sigma_{-i}(\theta_{-i}) \rangle, \theta_i) - u_i(\langle \sigma_i(\theta_i), \sigma_{-i}(\theta_{-i}) \rangle, \theta_i)]. \tag{1}$$

The *max regret* of strategy σ_i w.r.t. prior T, given θ_i and σ_{-i} is

$$MR_i(\sigma_i | \theta_i, T, \sigma_{-i}) = \max_{\theta_{-i} \in T(\theta_i)} R_i(\sigma_i | \theta_i, \theta_{-i}, \sigma_{-i}). \tag{2}$$

Finally, a *minimax best response* of agent i to σ_{-i} w.r.t. T is any strategy σ_i^* satisfying, for all $\theta_i \in \Theta_i$:

$$\sigma_i^* \in \arg\min_{\sigma_i \in \Sigma_i} MR_i(\sigma_i | \theta_i, T, \sigma_{-i}). \tag{3}$$

Intuitively, if we fix the behavior and types of all other agents, the regret of agent i with type θ_i for playing σ_i is the loss i experiences by playing σ_i rather than acting optimally. Of course, agent i does not know the true types of the other agents. The max regret of σ_i given prior T is the most i could regret playing σ_i (against the fixed strategies of the others) should an adversary choose its opponents's types in a manner consistent with its beliefs. Finally, a minimax best response is any strategy that minimizes this worst case loss in the face of such an adversary. Note that this strategy requires a minimax optimal choice for every possible type agent i could possess.

Unlike standard best responses, minimax best responses require agents to adopt a cautious stance with respect to possible realizations of opponent types. Without probabilistic information quantifying type uncertainty, minimax regret seems like the most natural decision criterion that could be adopted by such agents.

We define the notion of a *minimax-regret equilibrium* for a strict incomplete information game by analogy with Bayes-Nash equilibrium.

[2] Strict incomplete information games are equivalent to games in *informational form* [9].

Definition 2 A strategy profile σ is a *minimax-regret equilibrium* iff σ_i is a minimax best response to σ_{-i} for all agents i.

We note that other notions of qualitative equilibria have been proposed, but none have the same flavor as minimax-regret equilibria. Tennenholtz [12] describes qualitative equilibria for complete information games that rely on maximin strategies; but these do not have a clear extension to incomplete information games with type uncertainty. Work on *uncertainty aversion* can be viewed as incorporating some form of strict type uncertainty, but in a very different way. Rather than truly qualitative uncertainty, each agent is assumed to have a *set* of probabilistic priors (thus combining qualitative and quantitative uncertainty) [11]. Recently, equilibrium analysis of various auctions has been considered using this notion [4, 8]. Analysis of games in informational form naturally bears the closest relation to our work; however, to date, only *ex post* equilibria have been proposed for such games [9], which are considerably stronger than minimax-regret equilibria, and are not guaranteed to exist. In fact ex-post equilibria correspond to minimax-regret equilibria where the regret level of all the agents is zero.

Dominant strategies for strict incomplete information games can be defined in a similar way: we say σ_i is *minimax-dominant* if it is a minimax best response for *any* strategies σ_{-i} adopted by other players.[3] A *minimax-dominant strategy equilibrium (minimax-DSE)* is any strategy profile consisting of minimax-dominant strategies.

2.2 Existence of Equilibria

Not surprisingly, pure strategy minimax-regret equilibria for strict incomplete information games do not always exist (as is the case for BNE in Bayesian games). However, we can show that mixed-strategy minimax-regret equilibria exist for any finite game.

Theorem 1 *A mixed strategy minimax-regret equilibrium exists for any strict incomplete information game with finite agent, action, and type spaces.*

Proof sketch: The result can be proved using a similar strategy to classic proofs of the existence of (Bayes) Nash equilibria for finite games. We use Kakutani's fixed point theorem to show that the minimax-best-response correspondence (i.e., the mapping from any strategy profile to the set of profiles obtained by composing individual best responses to it) has a fixed point—this, by definition, is a minimax equilibrium. To apply the theorem, we show that minimax-best-response set for any profile is convex, and that the correspondence is upper-hemicontinuous. This relies on the piecewise-linear, convex nature of the max regret function itself (thus having a rather different character than best response correspondences based on expected utility). Moreover, minimax equilibria exist under most more general conditions under which Bayes-Nash equilibria exist with more complex type and action spaces (e.g., those identified in [7]).

Because dominant strategies in Bayesian games do not rely on the precise form on the prior, but only the set of possible types, we have (not surprisingly):

Proposition 1 *Strategy profile σ is a minimax-DSE for a strict incomplete information game with prior $T \subseteq \Theta$ iff it is a DSE for Bayesian game with type set T.*

[3] One could alternatively define dominant strategies in a "prior independent" way by requiring that regret be defined w.r.t. *any* type vector in Θ_{-i}. This would be somewhat more consistent with the typical definition in Bayesian games. This version will prove useful later.

3 Mechanism Design

Agents adopting minimax-regret equilibria use the minimax-regret decision criteria in a “standard” way, basically, by not quantifying their uncertainty over possible “states of nature” (i.e., the types of other agents). Since we typically view types as corresponding to utility functions, this is often very convenient. However, minimax-regret can also be exploited in mechanism design in a fashion that is more consistent with the recent use of the concept in preference elicitation.

In this context, the designer himself uses the minimax-regret criterion to deal with the (strict) uncertainty over which joint action the agents will take. Typically, the designer will be constrained to choose mechanisms that satisfy certain properties like individual rationality and incentive compatibility (in dominant strategies, minimax-regret or even Bayes-Nash equilibria). Minimax-regret is especially appropriate in the context of *partial revelation* mechanisms, where agents only reveal their true type partially, since in this case type uncertainty persists even after the agents have made their reports.

In [5], we explored the use of minimax-regret to optimize the automated design of partial revelation mechanisms. That is, we provide an efficient algorithm that, given a specific (strict) prior over types, outputs the partial revelation mechanism that minimizes the designer’s regret with respect to his objective function (e.g., social welfare, revenue). We are currently studying the use of minimax-regret in designing *sequential* mechanisms, where, as in preference elicitation, regret can help guide the partial revelation towards an optimal outcome.

References

[1] David E. Bell. Regret in decision making under uncertainty. *Operations Research*, 30:961–981, 1982.

[2] Allan Borodin and Ran El-Yaniv. *Online Computation and Competitive Analysis*. Cambridge Univ. Press, 1998.

[3] Craig Boutilier, Relu Patrascu, Pascal Poupart, and Dale Schuurmans. Constraint-based optimization with the minimax decision criterion. In *Proc. of CP-03*, pages 168–182, Kinsale, Ireland, 2003.

[4] Larry Epstein and JianJun Miao. A two-person dynamic equilibrium under ambiguity. *Journal of Economic Dynamics and Control*, 27(2):1253–1288, 2003.

[5] Nathanael Hyafil and Craig Boutilier. Regret minimizing equilibria and mechanisms for games with strict type uncertainty. In *Proc. of UAI-04*, pages 268–277, Banff, AB, Canada, 2004.

[6] Graham Loomes and Robert Sugden. Regret theory: An alternative theory of rational choice under uncertainty. *Economic Journal*, 92:805–824, 1982.

[7] Paul Milgrom and Rober J. Weber. Distributional strategies for games with incomplete information. *Mathematics of Operations Research*, 10(3):619–632, 1987.

[8] Emre Ozdenoren and Dan Levin. Auctions with a set of priors: Uncertain number of bidders. *Journal of Economic Theory*, 2004. Forthcoming.

[9] D. Monderer R. Holzman, N. Kfir-Dahav and Moshe Tennenholtz. Bundling equilibrium in combinatorial auctions. *Games and Economic Behavior*, page ???, 2003.

[10] Leonard J. Savage. *The Foundations of Statistics*. Wiley, New York, 1954.

[11] David Schmeidler and Itzhak Gilboa. Maxmin expected utility with a non-unique prior. *Journal of Mathematical Economics*, 18(2):141–154, 1989.

[12] Moshe Tennenholtz. On stable social laws and qualitative equilibria. *Artificial Intelligence*, 108:1–20, 1998.

[13] Tianhan Wang and Craig Boutilier. Incremental utility elicitation with the minimax regret decision criterion. In *Proc. of IJCAI-03*, pages 309–316, Acapulco, 2003.