Minimax Regret Treatment Choice with Incomplete Data and Many Treatments

Jörg Stoye*

New York University

January 27, 2006

Abstract

This note adds to the recent research project on treatment choice under ambiguity. I generalize Manski's (in press) analysis of minimax regret treatment choice by considering a more general setting and, more importantly, by solving for the treatment rule given finitely many (as opposed to two) treatments. The most interesting finding is that with three or more undominated treatments, the minimax regret treatment rule may assign the same treatment to all subjects; thus, the most salient feature of the two-treatment case does not generalize.

*I thank Chuck Manski, the co-editor, and especially an anonymous referee for helpful comments. Financial support through the Robert Eisner Memorial Fellowship, Department of Economics, Northwestern University, as well as a Dissertation Year Fellowship, The Graduate School, Northwestern University, is gratefully acknowledged. Address: Jörg Stoye, Department of Economics, New York University, 269 Mercer Street 7FL, New York, NY 10003, j.stoye@nyu.edu.

1 Motivation and Results

Consider a decision maker who must assign subjects to one of several treatments. The success of treatments is measured by some outcome variable Y, and the decision maker is an expected value maximizer, so if she knew the expected outcomes induced by different treatments, her problem would be trivial. However, her information is incomplete; specifically, she only knows that one of a set $\mathcal{S}$ of states of the world, identified with mappings from treatments to expected outcomes, has obtained. Since no probability distribution over $\mathcal{S}$ is provided, she faces a decision problem under ambiguity (or "Knightian uncertainty").

Manski (in press, see also Brock 2004) recently analyzed the special case of this problem that arises when the ambiguity is due to missing data, proposing *minimax regret* as the decision criterion and analyzing some of its features. This note extends his analysis by considering a more general decision environment and, more importantly, by allowing for finitely many as opposed to two treatments. The binary problem's most intriguing feature, namely that the minimax regret treatment rule assigns positive fractions of the population to either treatment unless one of the treatments is dominant, turns out not to generalize.

It is helpful to first restate the problem analyzed by Manski. Let there be treatments $t = 1, \dots, T$ that induce (random) potential outcomes $(Y_t)_{t=1}^T$ over a bounded support, say $Y_t \in [0, 1]$. If $\mu_t \equiv E(Y_t)$ were known for every t, treatment would be assigned accordingly. But observations of Y_t are plagued by missing data, and only a fraction p_t of realizations is observed. Defining $E_1(Y_t) \equiv E(Y_t|Y_t \text{ is observable})$ and $E_0(Y_t) \equiv E(Y_t|Y_t \text{ is missing})$, the decision maker's knowledge of μ_t is summarized by

$$\mu_t = p_t E_1(Y_t) + (1 - p_t) E_0(Y_t),$$

where $E_1(Y_t)$ and p_t are known but $E_0(Y_t)$ is not. (I will abstract from estimation problems, thus "observable" is synonymous with "known.") It follows that $\mu_t \in [p_t E_1(Y_t), p_t E_1(Y_t) + 1 - p_t]$, but this information may not suffice to identify the best treatment; nonetheless, some treatment rule has to be adopted. Manski (in press) proposes to evaluate treatment rules according to minimax regret. To formalize this, denote treatment rules by $\delta \in \Delta^{T-1}$, where Δ^{T-1} is the $(T-1)$-dimensional unit simplex in $\mathbb{R}^T$ and the t-th component of δ, δ_t henceforth, is the probability assigned to treatment t. Also letting $\mathcal{S} \equiv \times_t [p_t E_1(Y_t), p_t E_1(Y_t) + 1 - p_t]$ collect all possible configurations of $(\mu_t)_{t=1}^T$, the minimax regret ranking is given by

$$\delta \succeq \delta' \iff \max_{s \in \mathcal{S}} \left\{ \max_{d \in \Delta^{T-1}} \left\{ \sum_t d_t \mu_t \right\} - \sum_t \delta_t \mu_t \right\} \leq \max_{s \in \mathcal{S}} \left\{ \max_{d \in \Delta^{T-1}} \left\{ \sum_t d_t \mu_t \right\} - \sum_t \delta'_t \mu_t \right\}.$$

Since $\sum_t \delta_t \mu_t$ is the expected outcome induced by treatment rule δ, this can be thought of as a maximin criterion but with respect to ex-post efficiency loss rather than utility. For motivations and further discussions, see Manski (2004) and Stoye (2005b); the latter also contains many historical references.

Manski (in press) finds the minimax regret treatment rule when $T = 2$. A salient observation, replicated in Brock (2004) and as corollary 1 below, is that unless the better treatment is identified from the data, this treatment rule is mixed, i.e. it allocates a positive fraction of the population to each treatment.

I generalize this result in two ways. Firstly, I provide the solution for arbitrary finite T. Secondly, I allow for different informational settings as captured by information sets $\mathcal{S}$. Specifically, the only structure imposed on $\mathcal{S}$ is that if $\underline{\mu}_t \equiv \inf_{s \in \mathcal{S}} \{\mu_t\}$ and $\overline{\mu}_t \equiv \sup_{s \in \mathcal{S}} \{\mu_t\}$, then $\mathcal{S}$ contains all of the states $\{s_t\}_{t=1}^T$, where $s_t \equiv \left(\underline{\mu}_1, \underline{\mu}_2, \ldots, \underline{\mu}_{t-1}, \overline{\mu}_t, \underline{\mu}_{t+1}, \ldots, \underline{\mu}_T\right)$. Intuitively, this amounts to two restrictions:

- The bounds that one is able to place on μ_t are tight, i.e. they are achieved in some feasible state of the world s.
- Cross-restrictions between states are limited; in particular, it is possible that some treatment t is as good as possible, and all other treatments are as bad as possible, at the same time.

Manski's (in press) scenario emerges as special case with $\mathcal{S}$ as given above, but the requirement that $\{s_t\}_{t=1}^T \subseteq \mathcal{S}$ is significantly weaker. Bounds on expectations can arise from many other sources of partial identification as surveyed in Manski (2003), interval data and bounds on c.d.f.'s being but two examples. Similar informational settings also occur in the literature on interval probabilities (Walley 1991) and in Robust Bayesian inference (Wasserman & Kadane 1992). As long as there are no or appropriately limited cross-restrictions, the present analysis will apply to all of these cases.

The according generalization of proposition 1 in Manski (in press) is stated below. It is established by analyzing a fictitious zero-sum game between the decision maker and a malicious Nature whose expected utility equals the decision maker's expected regret. The Nash equilibria of this game are known to characterize all minimax regret treatment rules. In a first step, I show that Nature's action set can be reduced to $\{s_t\}_{t=1}^T$ (hence the requirement that $\mathcal{S}$ contain this set). It is then possible to analytically characterize the equilibria. A complication arises because there are two mutually exclusive types of them, captured below by cases (i) and (ii) with treatment rules δ respectively δ^*. (At their boundary in parameter space, a hybrid case obtains and is labeled (iii).) In either equilibrium, the expected utility of action s_t to Nature turns out to be indexed by a weighted average $\lambda_t \equiv \delta_t \underline{\mu}_t + (1-\delta_t)\overline{\mu}_t$ of the according treatment's best- and worst-case expected outcome, the weight being the probability with which the treatment is assigned. With this in mind, the expressions for δ respectively δ^* can be derived from Nature's best-response condition.

Proposition 1 ***Minimax Regret Treatment Rules for Arbitrary* T**

Order treatments such that $\overline{\mu}_1 \geq \ldots \geq \overline{\mu}_T$, where the tie-breaking rule is arbitrary. Let $\nu \equiv \max_t \left\{\underline{\mu}_t\right\}$, call a treatment t maximin if $\underline{\mu}_t = \nu$, and define $\delta \in \Delta^{T-1}$ and $\lambda \in \mathbb{R}$ implicitly by

$$\delta_t = \max\left\{\frac{\overline{\mu}_t - \lambda}{\overline{\mu}_t - \underline{\mu}_t}, 0\right\}, \forall t \in \{1, \ldots, T\}$$

and $t^* \in \{1,\ldots,T\}$ *by*

$$t^* \equiv \max\left\{N : \sum_{n=1}^{N} \frac{\nu - \underline{\mu}_n}{\overline{\mu}_n - \underline{\mu}_n} \leq 1\right\}$$

with the convention that $t^* \equiv T+1$ *if* $\sum_{t=1}^{T} \frac{\nu-\underline{\mu}_t}{\overline{\mu}_t-\underline{\mu}_t} < 1$*. Consider the following cases.*

(i) *If* $t^* > \max\{t : \delta_t > 0\}$*, then* $\delta^* = \delta$ *is the unique minimax regret treatment rule.*

(ii) *If* $t^* \leq \max\{t : \delta_t > 0\}$ *and* $\sum_{t=1}^{t^*} \frac{\nu-\underline{\mu}_t}{\overline{\mu}_t-\underline{\mu}_t} < 1$*, then* $\delta^* \in \Delta^{T-1}$ *is a minimax regret treatment rule iff*

$$\delta_t^* \begin{cases} \geq \max\left\{\frac{\overline{\mu}_t-\overline{\mu}_{t^*}}{\overline{\mu}_t-\underline{\mu}_t}, 0\right\}, & t \text{ is maximin} \\ = \max\left\{\frac{\overline{\mu}_t-\overline{\mu}_{t^*}}{\overline{\mu}_t-\underline{\mu}_t}, 0\right\}, & \text{otherwise} \end{cases}.$$

(iii) *Otherwise,* $\delta^* \in \Delta^{T-1}$ *is a minimax regret treatment rule iff*

$$\delta_t^* \begin{cases} \geq \max\left\{\frac{\overline{\mu}_t-\lambda^*}{\overline{\mu}_t-\underline{\mu}_t}, 0\right\}, & t \text{ is maximin} \\ = \max\left\{\frac{\overline{\mu}_t-\lambda^*}{\overline{\mu}_t-\underline{\mu}_t}, 0\right\}, & \text{otherwise} \end{cases}$$

for some $\lambda^* \in \left[\max\left\{\lambda, \overline{\mu}_{t^*+1}\right\}, \overline{\mu}_{t^*}\right]$*.*

A sufficient condition for the minimax regret treatment rule to be unique is that the maximin treatment is unique and $\sum_{t=1}^{t^*} \frac{\nu-\underline{\mu}_t}{\overline{\mu}_t-\underline{\mu}_t} \neq 1$*.*

Corollary 1 *Let* $T = 2$*, then the unique minimax regret treatment rule is*

$$\delta^* = \begin{cases} (0,1), & \underline{\mu}_2 \geq \overline{\mu}_1 \\ \left(\frac{\overline{\mu}_1-\underline{\mu}_2}{\overline{\mu}_1-\underline{\mu}_2+\overline{\mu}_2-\underline{\mu}_2}, \frac{\overline{\mu}_2-\underline{\mu}_1}{\overline{\mu}_1-\underline{\mu}_1+\overline{\mu}_2-\underline{\mu}_2}\right), & \min\{\overline{\mu}_1,\overline{\mu}_2\} > \max\{\underline{\mu}_1,\underline{\mu}_2\} \\ (1,0), & \underline{\mu}_1 \geq \overline{\mu}_2 \end{cases}.$$

2 Proofs and Discussion

Proposition 1. Consider the following zero-sum game between the decision maker and Nature: (i) The decision maker chooses a treatment rule $\delta \in \Delta^{T-1}$. Nature chooses a potentially mixed strategy $\sigma \in \Delta\mathcal{S}$, that is, a distribution over states of the world. (ii) A neutral meta-player draws $s \in \mathcal{S}$ according to σ. (iii) The decision maker's loss is $R(\delta, s) \equiv \max_t \{\mu_t\} - \sum_t \delta_t \mu_t$. Both players maximize expected utility. Then there exist convex sets $\mathcal{D}^* \subseteq \Delta^{T-1}$ and $\mathcal{S}^* \subseteq \Delta S$ s.t. (δ^*, σ^*) is a Nash equilibrium of the game iff $(\delta^*, \sigma^*) \in \mathcal{D}^* \times \mathcal{S}^*$; furthermore, $\mathcal{D}^*$ is the set of minimax regret treatment rules (e.g., see the appendix of Stoye 2005a).

Assume treatments are ordered as stated in the proposition. Fix any treatment rule δ. Any best response to δ is supported on

$$\arg\max_{s\in\mathcal{S}} R(\delta, s) = \arg\max_{s\in\mathcal{S}} \left\{\max_t \{\mu_t\} - \sum_t \delta_t \mu_t\right\}.$$

The objective function increases in μ_m iff $\mu_m = \max_t\{\mu_t\}$ and decreases in it otherwise (weakly so if $\delta_t = 0$). Thus, Nature always has a best response in the set $\{s_t\}_{t=1}^T$ as defined above. Consider the restricted game in which Nature's strategy space is reduced to this set; then it follows that any Nash equilibrium of this game is a Nash equilibrium of the unrestricted one. Since furthermore, every minimax regret treatment rule δ^* is a best response to every worst-case prior σ^*, an exact characterization of the minimax regret treatment rules for the restricted game implies an exact characterization for the unrestricted one. I will now provide the former.

In the restricted strategy space, nature's strategy σ can be identified with a vector $\sigma \in \Delta^{T-1}$, where $\sigma_t = \Pr(s = s_t)$. Observe that given σ, treatment t has expected outcome $\mu_t = \left(\sigma_t \overline{\mu}_t + (1-\sigma_t)\underline{\mu}_t\right)$.

Let (δ^*, σ^*) be a Nash equilibrium. To be a best response to σ^*, δ^* must solve

$$\begin{aligned} & \min_{\delta \in \Delta^{T-1}}! \left\{ E_{s|\sigma^*} \left(\max_t \{\mu_t\} - \sum_t \delta_t \mu_t \right) \right\} \\ = \quad & \min_{\delta \in \Delta^{T-1}}! \left\{ E_{s|\sigma^*} \left(\max_t \{\mu_t\} \right) - E_{s|\sigma^*} \left(\sum_t \delta_t \mu_t \right) \right\}. \end{aligned}$$

The left-hand expectation does not vary with δ, so that δ^* must maximize $E_{s|\sigma^*}(\Sigma_t \delta_t \mu_t)$, i.e. expected outcome given σ^*. Hence, $\delta_t^* > 0$ only if $\sigma_t^* \overline{\mu}_t + (1-\sigma_t^*)\underline{\mu}_t = \mu^* \equiv \max_t \left\{\sigma_t^* \overline{\mu}_t + (1-\sigma_t^*)\underline{\mu}_t\right\}$. Furthermore, (δ^*, σ^*) is known to induce expected outcome μ^*; observe in particular that $\mu^* \geq \nu$.

If σ^* maximizes regret against δ^*, it must also maximize regret against δ^* under the additional constraint that $\sigma_t^* \overline{\mu}_t + (1-\sigma_t^*)\underline{\mu}_t \leq \mu^*, \forall t$, because by construction, this constraint will not bind at σ^*. Hence, σ^* must solve

$$\begin{aligned} & \max_{\sigma \in \Delta^{T-1}}! \left\{ \sum_t \sigma_t \overline{\mu}_t - \mu^* \right\} \\ s.t. \quad & \sigma_t^* \overline{\mu}_t + (1-\sigma_t^*)\underline{\mu}_t \leq \mu^*, \forall t, \end{aligned}$$

where I also use the knowledge that in equilibrium, $E_{s|\sigma^*}\left(\sum_t \delta_t \mu_t\right) = \mu^*$.

The above problem is linear in every $\overline{\mu}_t$; since treatments are ordered according to $\overline{\mu}_t$, it is solved by maximizing the probability mass placed on lower-numbered treatments:

$$\sigma_t^* = \begin{cases} \frac{\mu^* - \underline{\mu}_t}{\overline{\mu}_t - \underline{\mu}_t}, & t < \widetilde{t} \\ 1 - \sum_{t \neq \widetilde{t}} \sigma_t^*, & t = \widetilde{t} \\ 0, & t > \widetilde{t} \end{cases}, \qquad (1)$$

where $\widetilde{t} \equiv \max\left\{N \leq T : \Sigma_{n=1}^N \frac{\mu^* - \underline{\mu}_t}{\overline{\mu}_t - \underline{\mu}_t} \leq 1\right\}$.

But σ^* being a best response to δ^* also requires that

$$\begin{aligned} \forall t, t' \text{ s.t. } \sigma_t^*, \sigma_{t'}^* > 0 : R(\delta^*, s_t) &= R(\delta^*, s_{t'}) \\ \Longrightarrow \sum_{i \neq t} \delta_i^* (\overline{\mu}_t - \underline{\mu}_i) &= \sum_{i \neq t'} \delta_i^* (\overline{\mu}_t - \underline{\mu}_i) \\ \Longrightarrow (1-\delta_t^*)\overline{\mu}_t - \delta_{t'}^* \underline{\mu}_{t'} &= (1-\delta_{t'}^*)\overline{\mu}_{t'} - \delta_t^* \underline{\mu}_t \\ \Longrightarrow \delta_t^* \underline{\mu}_t + (1-\delta_t^*)\overline{\mu}_t &= \delta_{t'}^* \underline{\mu}_{t'} + (1-\delta_{t'}^*)\overline{\mu}_{t'} \end{aligned}$$

and

$$\begin{aligned}\forall t,t' \text{ s.t. } \sigma_t^* > 0 = \sigma_{t'}^* : R(\delta^*, s_t) &\geq R(\delta^*, s_{t'}) \\ \Longrightarrow \delta_t^* \underline{\mu}_t + (1-\delta_t^*)\overline{\mu}_t &\geq \delta_{t'}^* \underline{\mu}_{t'} + (1-\delta_{t'}^*)\overline{\mu}_{t'}.\end{aligned}$$

Hence, δ^* is partially characterized by the fact that there exists λ s.t.

$$\begin{aligned}\delta_t^* \underline{\mu}_t + (1-\delta_t^*)\overline{\mu}_t &\quad \begin{cases} = \lambda^*, & \sigma_t^* > 0 \\ \leq \lambda^*, & \sigma_t^* = 0 \end{cases} \\ \Longleftrightarrow \delta_t^* &\quad \begin{cases} = \frac{\overline{\mu}_t - \lambda^*}{\overline{\mu}_t - \underline{\mu}_t}, & \sigma_t^* > 0 \\ \geq \frac{\overline{\mu}_t - \lambda^*}{\overline{\mu}_t - \underline{\mu}_t}, & \sigma_t^* = 0 \end{cases} .\end{aligned} \tag{2}$$

Of course, δ_t^* must also be nonnegative, so that $\frac{\overline{\mu}_t - \lambda^*}{\overline{\mu}_t - \underline{\mu}_t}$ can be replaced with $\max\left\{\frac{\overline{\mu}_t - \lambda^*}{\overline{\mu}_t - \underline{\mu}_t}, 0\right\}$ in [2]. To proceed with the characterization, one must distinguish two cases.

Case 1: For all t, $\delta_t^* = \max\left\{\frac{\overline{\mu}_t - \lambda^*}{\overline{\mu}_t - \underline{\mu}_t}, 0\right\}$. Since also $\Sigma_t \delta_t^* = 1$, this case generates $(T+1)$ linear equations in $(T+1)$ unknowns. If one guesses $\lambda^* = \max_t \{\overline{\mu}_t\}$, then it is easily seen that $\Sigma_t \delta_t^* = 0$ is implied. On the other hand, a guess of $\lambda^* = \underline{\mu}_1$ implies that $\delta_1^* = 1$ and hence $\Sigma_t \delta_t^* \geq 1$. For intermediate values of λ^*, the implied value of $\Sigma_t \delta_t^*$ is continuously decreasing in λ^*, hence $\Sigma_t \delta_t^* = 1$ holds for exactly one λ. The case therefore generates a unique candidate δ^*.

Best response conditions now constrain σ^* as follows:

$$\begin{aligned}\delta_t^* &> 0 \Longrightarrow \sigma_t^* \overline{\mu}_t + (1-\sigma_t^*)\underline{\mu}_t = \mu^* \Longrightarrow \sigma_t^* = \frac{\mu^* - \underline{\mu}_t}{\overline{\mu}_t - \underline{\mu}_t} \\ \delta_t^* &= 0 \Longrightarrow \sigma_t^* \overline{\mu}_t + (1-\sigma_t^*)\underline{\mu}_t \leq \mu^* \Longrightarrow \sigma_t^* \leq \frac{\mu^* - \underline{\mu}_t}{\overline{\mu}_t - \underline{\mu}_t}.\end{aligned}$$

If $t^* > \max\{t : \delta_t > 0\}$, then the above constraints are consistent with [1]. It can then be straightforwardly verified that (δ^*, σ^*) constitutes a Nash equilibrium and that δ^* is characterized as in case (i). If $t^* \leq \max\{t : \delta_t > 0\}$ and the above constraints imply that $\Sigma_t \sigma_t^* > 1$, the equilibrium fails. If $\Sigma_t \sigma_t^* \geq 1$ is implied, the equilibrium will exist but may be nested within case (iii).

Case 2: For some t, $\delta_t^* > \max\left\{\frac{\overline{\mu}_t - \lambda}{\overline{\mu}_t - \underline{\mu}_t}, 0\right\}$. This implies that the decision maker chooses with positive probability some treatments $t', t'', \ldots$ with $\sigma_{t'}^* = \sigma_{t''}^* = \ldots = 0$, hence their expected outcomes are equal to $\underline{\mu}_{t'}, \underline{\mu}_{t''}, \ldots$. This is consistent with δ^* being a best response only if $\underline{\mu}_{t'} = \underline{\mu}_{t''} = \ldots = \mu^*$; since $\mu^* \geq \nu$, it must be the case that all of these treatments are maximin and that $\mu^* = \nu$.

Expression [1] now fully characterizes σ^*. If $\sum_{t=1}^{t^*} \frac{\nu - \underline{\mu}_t}{\overline{\mu}_t - \underline{\mu}_t} > 1$, then

$$\sigma_{t^*}^* < \frac{\nu - \underline{\mu}_t}{\overline{\mu}_t - \underline{\mu}_t} \Longrightarrow \delta_{t^*}^* = 0.$$

This is consistent with [2] only if $\overline{\mu}_{t^*} = \lambda^*$. With λ^* thus defined, expression [2] bounds every δ_t^* from below. If these lower bounds sum to more than one, the candidate equilibrium fails. Otherwise, every δ^* that fulfils the bounds and has $\Sigma_t \delta_t^* = 1$ is a best response to σ^*. (If there is exactly one maximin treatment, δ^* is unique.)

In the knife-edge case that $\sum_{t=1}^{t^*} \frac{\nu - \underline{\mu}_t}{\overline{\mu}_t - \underline{\mu}_t} = 1$, λ^* may not be uniquely determined. However, inspection of [2] reveals that $\lambda^* \leq \overline{\mu}_{t^*}$ (because $\delta_{t^*}^* \geq 0$) and that $\lambda^* \geq \overline{\mu}_{t^*+1}$ (because $\delta_{t^*+1}^* = 0$). An alternative lower bound on λ^* is generated by assuming that all inequalities in [2] bind, i.e. by presuming case (i). (If this constraint binds, both cases yield valid equilibria, with the set of minimax regret rules characterized under case 2 containing the δ^* characterized in case 1.) Any choice of λ^* that is consistent with these constraints will generate a δ^* s.t. (δ^*, σ^*) constitutes a Nash equilibrium.

That one of the cases will yield a Nash equilibrium is not obvious from the above. Notice, however, that the cases are exhaustive, and by existence of a Nash equilibrium, one of them must yield at least one equilibrium. If both cases yielded valid sets of equilibria with the equilibrium from case (i) not contained in the equilibria from case (ii), the set of Nash equilibria would be disconnected, contradicting convexity. Also, sufficiency of the conditions is easily established.■

Corollary 1. Let $T = 2$ and $\underline{\mu}_1 < \overline{\mu}_2$ as well as $\underline{\mu}_2 < \overline{\mu}_1$. (For the other cases, δ^* is obvious.) Assume by contradiction that $\mu^* = \nu$, then σ_t^* will be zero if t is maximin and will be less than one otherwise, implying that $\Sigma_t \sigma_t^* < 1$. Hence, case (i) applies, and [2] can be solved for δ^*.■

Mathematically, proposition 1 implies that the set of minimax regret treatment rules is the convex hull of a finite number of extremal treatment rules. Whilst somewhat tedious to describe, these extrema are easy to evaluate, and once identified, their validity can be checked by paper-and-pencil methods.[1] Indeed, the only generic condition under which the set is not a singleton is when case (ii) obtains and there are several maximin utility treatments, in which case only their aggregate probability will be determined. Uniqueness of δ^* will also fail in case (iii), but this case is nongeneric in the sense that it requires a linear equality in parameters to hold; intuitively, it connects the first two cases at their boundary in parameter space.[2]

Other than this, general intuitions about δ^* seem to be limited to the fact that it is always supported on the t^* best treatments in terms of the max max-ordering, i.e. according to $\overline{\mu}_t$, where t^* is some number between 1 and T. Two more features are illustrated by the following examples.

Example 1 *Let $T = 2$ and let bounds on μ_t be as follows:*

[1] MATLAB code that evaluates proposition 1 is posted on the author's webpage at `http://homepages.nyu.edu/~js3909/`.

[2] Examples of either form of non-unique solutions are given in an online appendix on the aforementioned webpage.

t	**1**	**2**
$\overline{\boldsymbol{\mu}}_t$	5	3
$\underline{\boldsymbol{\mu}}_t$	2	0

Then $\delta^* = (5/6, 1/6)$ *uniquely.*

Example 2 *Let* $T = 3$ *and let bounds on* μ_t *be as follows:*

t	**1**	**2**	**3**
$\overline{\boldsymbol{\mu}}_t$	5	3	3
$\underline{\boldsymbol{\mu}}_t$	2	0	0

Then $\delta^* = (1, 0, 0)$ *uniquely.*

Firstly, example 2 is generated from example 1 by adding a treatment that is seemingly like treatment 2.[3] As a result, preferences regarding the first two treatments change, and the second treatment, which was previously selected with positive probability, is not chosen any more. Some readers may find this counterintuitive – it is certainly not obvious that treatment 2 got worse due to the presence of treatment 3. The observation illustrates that minimax regret violates the well-known Independence of Irrelevant Alternatives axiom.

Secondly, δ^* in example 2 is not mixed, i.e. the most salient aspect of the two-treatment case does not carry over. This finding is of general interest because researchers investigating minimax regret decision rules usually find these to mix, often in contrast to maximin utility rules (Bergemann & Schlag 2005, Brock 2004, Manski 2004, Schlag 2003, Stoye 2005a). To understand why it obtains, it is instructive to informally prove that δ^* will always mix in the case of two (non-dominating) treatments. Suppose by contradiction that δ^* is concentrated at $t = 1$, then Nature's best response in the fictitious game is to deterministically choose state s_2. But the decision maker's best response to that would be to assign all subjects to treatment 2. The best response correspondence therefore cycles over pure strategies. In short, the result obtains because 2×2 anti-coordination games have only fully mixed equilibria.[4]

But this observation does not generalize to more complex games. In example 2 above, one of Nature's best responses to δ^* is to mix equally over s_2 and s_3; δ^* then remains the decision maker's strict best response and hence is the unique minimax regret treatment rule. Notice that $\overline{\mu}_2 = \overline{\mu}_3$ is necessary for this to work, thus a non-mixing δ^* obtains only in specific, although interesting, cases.

References

[1] Bergemann, D. & K.H. Schlag (2005) Robust Monopoly Pricing: The Case of Regret. Discussion paper, Yale University and European University Institute.

[3] The result depends, however, on the presumption that treatments 2 and 3 are not actually the same, because the possibility of their outcomes being different is crucial.

[4] In example 1 above, the equilibrium is completed by setting $\sigma^* = (1/6, 5/6)$.

[2] Brock, W.A. (2004) Profiling Problems with Partially Identified Structure. Discussion paper, University of Wisconsin-Madison.

[3] Manski, C.F. (2003) *Partial Identification of Probability Distributions.* Springer Verlag.

[4] — (2004) Statistical treatment rules for heterogeneous populations. *Econometrica* 72, 1221–1246.

[5] — (in press) Minimax-regret treatment choice with missing outcome data. *Journal of Econometrics.*

[6] Schlag, K.H. (2003) How to Minimize Maximum Regret under Repeated Decision-Making. Discussion paper, European University Institute.

[7] Stoye, J. (2005a) Minimax Regret Treatment Choice with Finite Samples. Discussion paper, New York University.

[8] — (2005b) Statistical Decisions under Ambiguity. Discussion paper, New York University.

[9] Walley, P. (1991) *Statistical Reasoning with Imprecise Probabilities.* Chapman and Hall.

[10] Wasserman, L. & J.B. Kadane (1992) Computing bounds on expectations. *Journal of the American Statistical Association* 87, 516-522.