

Minimax and the Value of Information*

Evan Sadler[†]

April 4, 2014

Abstract

In his discussion of minimax decision rules, Savage (1954, p. 170) presents an example purporting to show that minimax applied to negative expected utility (referred to by Savage as “negative income”) is an inadequate decision criterion for statistics; he suggests the application of a minimax regret rule instead. The crux of Savage’s objection is the possibility that a decision maker would choose to ignore even “extensive” information. More recently, Parmigiani (1992) has suggested that minimax regret suffers from the same flaw. He demonstrates the existence of “relevant” experiments that a minimax regret agent would never pay a positive cost to observe. On closer inspection, I find that minimax regret is more resilient to this critique than would first appear. In particular, there are cases in which no experiment has any value to an agent employing the minimax negative income rule, while we may always devise a hypothetical experiment that a minimax regret agent would pay for. The force of Parmigiani’s critique is further blunted by the observation that “relevant” experiments exist for which a Bayesian agent would never pay. I conclude with a discussion of pessimism in the context of minimax decision rules.

1 Introduction

In recent years there has been a revival of interest in the application of minimax procedures to statistical and econometric problems. Manski (2004) for instance, has spawned a literature on the application of minimax regret to treatment choice, and Stoye (2011) has carefully examined the axiomatic distinctions between different minimax rules. One reason for this newfound popularity is surely the increase in computing power now available to researchers, rendering such methods feasible in more domains. With renewed interest has come increased scrutiny. In particular, concerns over the use of information by minimax procedures have resurfaced. It has long been known that minimax applied to negative

*I am grateful to Roy Radner for calling my attention to this problem as well as for numerous fruitful discussions. I also thank Jörg Stoye for pointing me to useful references. Any errors are mine alone.

[†]Stern School of Business, New York University – esadler@stern.nyu.edu

expected utility (hereafter, “negative income”) can entirely ignore available information in some decision problems. Savage (1954, p. 170) presented a simple example demonstrating this phenomenon, asserting that minimax applied to negative income was entirely inadequate as a criterion for statistics. His alternative was applying minimax to regret. While many credit Savage with the invention of the minimax regret criterion, Savage himself gave priority to Wald (1950), believing that Wald could not possibly have intended minimax negative income in his work. Savage clearly believed this critique inapplicable to regret, but more recent work has disputed this.

Parmigiani (1992) analyzes an example in which minimax applied to negative income assigns positive value to an experiment, while minimax regret does not. He goes on to characterize the circumstances under which minimax rules refrain from “relevant” experimentation, even if the experiment is nearly costless. Parmigiani describes this phenomenon as *ultrapessimism*. The essence of his result is best understood if we interpret the statistician’s problem as a zero sum game played against nature. If nature has an equilibrium strategy that is supported on relatively few states, then experiments that fail to distinguish these states will be worthless. If such cases can be expected to occur regularly outside of contrived examples, then this raises serious concerns about the use of minimax regret in applied work.

One goal of the present paper is to controvert this criticism of minimax applied to regret. Savage’s example uncovered a more fundamental flaw in minimax negative income than the one Parmigiani discusses. In some decision problems, minimax negative income not only fails to utilize an observation that another procedure finds valuable, but in fact no conceivable experiment provides any value to an agent following a minimax negative income rule. This is in stark contrast to the situation we find with minimax regret, where we can always construct a valuable experiment. Moreover, even when valuable experiments do exist for the minimax negative income rule, there may still be some states for which no amount of evidence in their favor will cause the decision maker to adopt an action better suited to those states. Again, minimax applied to regret avoids such issues. On further examination, the characterization of *ultrapessimism* Parmigiani gives can apply to Bayesian procedures just as easily—even when a minimax regret agent would pay for the information. *Ultrapessimism*, so defined, fails to distinguish meaningfully between decision procedures, but Savage’s observation remains a substantive distinction between minimax regret and minimax negative income.

A second goal is to critically examine other objections to minimax regret. There are other senses in which we can describe the procedure as “pessimistic,” and these may present challenges to researchers in applications. Many issues stem from prescriptions that go against what we intuitively expect from a “good” estimator. In a sense, regret is criticized for failing to conform to our prior information, but a major motivation for minimax procedures is to avoid subjective priors. An important direction for future work lies in devising ways to utilize such information alongside minimax procedures.

In the next section I provide a detailed analysis of the examples given by Savage (1954) and Parmigiani (1992). The following section generalizes previous results that characterize when minimax values information. This section also establishes the key distinction between

minimax applied to negative income and minimax applied to regret. Finally, I offer some thoughts on the problem of “pessimism” in minimax estimation and the effective application of minimax regret.

2 Examples

In this section I present two examples; first, Savage’s (1954) example, followed by Parmigiani’s (1992). Consider a state space partitioned into two events, B_1 and B_2 , with two primary acts, f_1 and f_2 . Suppose f_1 yields a payoff of -1 utiles in both events, and f_2 yields -10 if B_1 obtains and 1 if B_2 obtains.

	B_1	B_2
f_1	-1	-1
f_2	-10	1

Assuming the decision maker is capable of making randomized choices, there are derived acts of the form $f = \lambda f_1 + (1 - \lambda)f_2$ for $\lambda \in [0, 1]$. The negative income from a particular act in each event is then,

$$-I(f; B_1) = 10 - 9\lambda, \quad -I(f; B_2) = 2\lambda - 1.$$

Note for all values of λ in the given range, $-I(f; B_1)$ is larger than $-I(f; B_2)$, so if our decision maker applies a minimax negative income decision rule, the chosen act will be the one that minimizes $-I(f; B_1)$, which is to simply choose f_1 with probability 1.

Now suppose the decision maker has the option to observe at no cost the random variable X taking values in $\{1, 2\}$ with

$$\begin{aligned} \mathbb{P}(X = 1 | B_1) &= \mathbb{P}(X = 2 | B_2) = 1 - g, \\ \mathbb{P}(X = 2 | B_1) &= \mathbb{P}(X = 1 | B_2) = g, \end{aligned}$$

for some $g < \frac{1}{2}$. There are now four primary acts as functions of X : $g_1(X) = f_1$ for each value of X , $g_2(X) = f_2$ for each value of X , g_3 defined by $g_3(1) = f_1$ and $g_3(2) = f_2$, and g_4 defined by $g_4(1) = f_2$ and $g_4(2) = f_1$. The set of derived acts available to the decision maker are the convex combinations of these four: $f = \lambda_1 g_1 + \lambda_2 g_2 + \lambda_3 g_3 + \lambda_4 g_4$, where $\lambda_1 + \lambda_2 + \lambda_3 + \lambda_4 = 1$. For such an act, we can calculate the negative income in each event as,

$$\begin{aligned} -I(f; B_1) &= \lambda_1 + 10\lambda_2 + (1 + 9g)\lambda_3 + (10 - 9g)\lambda_4, \\ -I(f; B_2) &= \lambda_1 - \lambda_2 + (2g - 1)\lambda_3 + (1 - 2g)\lambda_4. \end{aligned}$$

Just as before, we have $-I(f; B_1) \geq -I(f; B_2)$ for all feasible parameter values. In the expression for $-I(f; B_1)$, the coefficients of the λ_2 , λ_3 , and λ_4 terms are all strictly larger than 1, so this is minimized by allocating all weight to λ_1 . That is, the decision maker

chooses f_1 regardless of the observed X and regardless of the value g , earning utility of -1 in all cases. If observing X were costly, the agent would evidently decline the observation, no matter how small the cost. This is the basis for Savage's criticism of the minimax negative income rule.

Compare this result with that obtained using a minimax regret rule when faced with the same situation. Let $L(f; B_i)$ denote the regret of action f if B_i obtains. For the original problem in the absence of the observation, we have

$$L(f; B_1) = 9(1 - \lambda), \quad L(f; B_2) = 2\lambda.$$

The decision maker chooses λ to minimize the maximum of $L(f; B_1)$ and $L(f; B_2)$. Note $L(f; B_1)$ is strictly decreasing in λ , and $L(f; B_2)$ is strictly increasing. Therefore, the minimal maximum occurs when $L(f; B_1) = L(f; B_2)$, or $9(1 - \lambda) = 2\lambda$, implying $\lambda = \frac{9}{11}$. Substituting $\lambda = \frac{9}{11}$ into the equation for $L(f; B_2)$ yields the minimal regret

$$L(f; B_2) = 2\lambda = 2 \cdot \frac{9}{11} = \frac{18}{11}.$$

Given the same random variable X as before, and the same four primary acts, the situation changes. The regret we compute in each event is

$$L(f; B_1) = 9\lambda_2 + 9g\lambda_3 + 9(1 - g)\lambda_4, \quad L(f; B_2) = 2\lambda_1 + 2g\lambda_3 + 2(1 - g)\lambda_4.$$

Note since $g < \frac{1}{2}$, the coefficient of λ_3 is lower than that of λ_4 in both of these. Therefore, the minimal regret must occur with $\lambda_4 = 0$.

Now consider the following observations. If $\lambda_1 < (9\lambda_2 + 7g\lambda_3)/2$, then $L(f; B_1) > L(f; B_2)$, and we can reduce the maximum of $L(f; B_1)$ and $L(f; B_2)$ by reallocating a small amount of weight from either λ_2 or λ_3 to λ_1 . Therefore, we must have $\lambda_1 \geq (9\lambda_2 + 7g\lambda_3)/2$ and $L(f; B_1) \leq L(f; B_2)$. In fact, we must have equality. If $L(f; B_1) < L(f; B_2)$, then the maximum is reduced by reallocating some weight from λ_1 to λ_2 .

Write $\lambda_3 = 1 - \lambda_1 - \lambda_2$, and the condition $\lambda_1 = (9\lambda_2 + 7g\lambda_3)/2$ becomes

$$(9 - 9g)\lambda_2 - 9g\lambda_1 + 9g = (2 - 2g)\lambda_1 - 2g\lambda_2 + 2g \implies \lambda_1 = \frac{(9 - 7g)\lambda_2 + 7g}{2 + 7g}.$$

The regret is then

$$(9 - 9g)\lambda_2 - 9g \frac{(9 - 7g)\lambda_2 + 7g}{2 + 7g} + 9g = \frac{18 - 36g}{2 + 7g} \lambda_2 + 9g + \frac{63g^2}{2 + 7g}.$$

Since $g < \frac{1}{2}$, the coefficient of λ_2 is positive, and the regret is minimized at $\lambda_2 = 0$. Thus, the minimax regret act sets $\lambda_1 = 7g/(7g + 2)$, $\lambda_2 = 0$, and $\lambda_3 = 2/(7g + 2)$. We compute the minimal regret as

$$L(f; B_1) = 9g\lambda_3 = 9g \cdot \frac{2}{7g + 2} = \frac{18g}{7g + 2},$$

which goes to zero as g goes to zero (i.e. as X becomes more informative). The agent should be willing to pay as much as $\frac{18}{11} - 18g/(7g + 2) > 0$ to observe X .

For the second example, consider a box containing two ordered marbles in one of three arrangements. Either both are red, the first is blue with the second red, or both are blue. Two primary acts are available. The act f_1 pays -2 utiles if both are red, 0 if the first is blue and the second red, and -4 if both are blue. The act f_2 pays -4 if both are red or the first blue and the second red, and it pays 0 if both are blue. The payoff table is depicted below.

	RR	BR	BB
f_1	-2	0	-4
f_2	-4	-4	0

Considering mixed actions as before, the negative income from choosing $f = \lambda f_1 + (1 - \lambda)f_2$ in each state is

$$-I(f; RR) = 4 - 2\lambda, \quad -I(f; BR) = 4 - 4\lambda, \quad -I(f; BB) = 4\lambda.$$

We have $-I(f; BR) \leq -I(f; RR)$ for all $\lambda \in [0, 1]$, so we can solve for the minimax solution by taking $-I(f; RR) = -I(f; BB)$, implying $\lambda = \frac{2}{3}$. Substituting this value into the equation for $-I(f; BB)$ gives the expected negative income of $\frac{8}{3}$.

If the decision maker can observe the first marble at no cost, there are four primary acts as a function of the observation. In fact, we need only consider two of these acts since f_1 dominates f_2 whenever the observed marble is red. If the first marble is blue, then clearly the minimax solution is to choose f_1 and f_2 with equal probability, leading to expected negative income of 2 in each state. The decision maker should be willing to pay as much as $\frac{2}{3}$ to observe the first marble.

Now consider a decision maker using minimax regret in the same situation. The regret table is

	RR	BR	BB
f_1	0	0	4
f_2	2	4	0

The regret from the act $f = \lambda f_1 + (1 - \lambda)f_2$ in each state is

$$L(f; RR) = 2 - 2\lambda, \quad L(f; BR) = 4 - 4\lambda, \quad L(f; BB) = 4\lambda.$$

Since $L(f; BR) \geq L(f; RR)$ for all $\lambda \in [0, 1]$, we can solve for the minimax strategy by taking $L(f; BR) = L(f; BB)$, implying $\lambda = \frac{1}{2}$. The minimax regret strategy is to randomize equally between the two acts, yielding a maximum expected regret of 2 .

Given the opportunity to observe the first marble, a cursory examination of the payoff table shows that zero regret can be achieved in state RR , but will remain unchanged in the other two. The optimal strategy will be to randomize equally between the act that chooses f_1 always and the act that chooses f_1 only when the observed marble is red. The

maximum regret is still 2, so the decision maker should be unwilling to pay any positive amount for the observation based on the minimax regret criterion. Parmigiani (1992) claims that this example shows minimax regret is vulnerable to the same criticism as minimax negative income. In the next section, I shall rebut this claim while providing some insight into when these two decision rules will ignore available information.

3 The Value of an Experiment

We can understand the phenomenon demonstrated through these examples via the correspondence between the decision maker's problem and a two player zero sum game played between nature and the decision maker (Wald, 1950). Consider a decision problem with a finite collection of states Ω and a finite number of primary acts. Imagine the corresponding payoff table, either a negative income or a regret table. If we consider this problem as a zero sum game in which nature earns the negative of the decision maker's payoff, then it is well-known that the Nash equilibria of the game share a common value, and this is the value the decision maker will achieve through application of a minimax decision rule.

Let $\sigma = (\sigma_d, \sigma_n)$ denote an equilibrium strategy profile in the game described above, and let Ω_σ denote the collection of states that nature plays with positive probability under σ . Consider a random variable X taking values in some arbitrary measure space $(\mathcal{S}, \mathcal{S})$; we shall call such a random variable an experiment if the decision maker can observe its value prior to acting. We obtain the following result.

Proposition 1. *In the decision problem described above, the experiment X has no value to a minimax decision maker if there exists an equilibrium strategy profile σ such that*

$$\mathbb{P}(X \in S | \omega_1) = \mathbb{P}(X \in S | \omega_2)$$

for all $\omega_1, \omega_2 \in \Omega_\sigma$ and for all $S \in \mathcal{S}$.

If there are only two primary acts, and the experiment X has no value, then such an equilibrium strategy profile necessarily exists.

Proof. Observe first that the value of the game to the decision maker can be no lower with the experiment since the decision maker may always choose to ignore it. If we have $\mathbb{P}(X \in S | \omega_1) = \mathbb{P}(X \in S | \omega_2)$ for all $\omega_1, \omega_2 \in \Omega_\sigma$ and all $S \in \mathcal{S}$, then nature may play the strategy specified by σ and be sure to obtain the same value from the game with or without the experiment X . The value of the game is therefore unchanged; hence, the attainable minimax value is unchanged with the introduction of the experiment.

Conversely, suppose there are two primary acts, and after introducing the experiment, nature has a strategy σ_n that can ensure the same value of the game as without the experiment. Consider a minimal such strategy, in the sense that nature cannot ensure this payoff through randomization over a strict subset of the states played with positive probability in σ_n . Nature could have played σ_n in the game without the experiment, obtaining a value at

least as high; therefore, this strategy is part of an equilibrium strategy profile (σ_d, σ_n) in the game without the experiment.

If there exist $\omega_1, \omega_2 \in \Omega_\sigma$ and $S \in \mathcal{S}$ with

$$p_1 = \mathbb{P}(X \in S | \omega_1) \neq \mathbb{P}(X \in S | \omega_2) = p_2,$$

I claim that the decision maker can obtain a strictly higher payoff if nature plays σ_n after the introduction of the experiment. To see why, consider two cases. First suppose that σ_n is minimal in the game without the experiment. Given there are two primary acts f_1 and f_2 , we can relabel ω_1 and ω_2 such that the payoff from f_1 is at least as high as that from f_2 in state ω_1 , and the payoff from f_2 is at least as high as that from f_1 in state ω_2 . One of these inequalities is strict: otherwise σ_n is clearly not minimal (we could remove either ω_1 or ω_2 from the support of nature's strategy). We can partition Ω_σ into Ω_1 and Ω_2 , with $\omega_1 \in \Omega_1$, $\omega_2 \in \Omega_2$, and

$$\mathbb{P}(X \in S | \omega \in \Omega_1) \neq \mathbb{P}(X \in S | \omega \in \Omega_2).$$

Without loss of generality, suppose the first probability is greater. Consider the strategy $\sigma_d^{(X)}$, which plays f_1 if $X \in S$ and f_2 if $X \in S^c$, and note that for small enough λ the strategy $\lambda\sigma_d^{(X)} + (1 - \lambda)\sigma_d$ will obtain a strictly higher payoff than σ_d in response to σ_n .

Alternatively, if σ_n is no longer minimal when the experiment is removed from the game, we can find a minimal equilibrium strategy σ'_n for nature to follow in the game without the experiment, which is supported on a strict subset of the states supporting σ_n . By the argument in the preceding paragraph, we cannot have $\omega_1, \omega_2 \in \Omega_{\sigma'}$ and $S \in \mathcal{S}$ with

$$p_1 = \mathbb{P}(X \in S | \omega_1) \neq \mathbb{P}(X \in S | \omega_2) = p_2.$$

Otherwise the experiment allows the decision maker to make a strict improvement. However, the argument in the first paragraph of this proof implies that σ'_n is an equilibrium strategy of nature in the game with the experiment, guaranteeing the same payoff. This contradicts the minimality of σ_n in the game with the experiment. Therefore, σ_n is minimal in the game without the experiment, and we have

$$\mathbb{P}(X \in S | \omega \in \Omega_1) \neq \mathbb{P}(X \in S | \omega \in \Omega_2)$$

for all $\omega_1, \omega_2 \in \Omega_\sigma$ and $S \in \mathcal{S}$. □

I note in passing that the restriction to two primary acts in the second part of the proposition is needed to rule out the situation depicted in the following regret table, in which many (but not all!) informative experiments are insufficient to cause a minimax agent to deviate from f_3 .

	ω_1	ω_2
f_1	0	10
f_2	10	0
f_3	1	1

The key insight of this result is that an experiment is worthless to a minimax agent if it does not discriminate between the states that nature plays in equilibrium. Put this way, it sounds rather unremarkable, but it explains the unwillingness of minimax agents to pay for information in the examples given previously. If nature makes use of only a small subset of available states in equilibrium, then the minimax agent's decision process is entirely dominated by this subset of states, and experiments giving information about other states will have no value. In Savage's example, nature has a dominant strategy in the negative income game, and consequently there can never be an experiment with value to a minimax negative income agent. In Parmigiani's example, state RR is weakly dominated in the regret game, but not in the negative income game.

Parmigiani (1992) proves a similar result under somewhat stronger conditions, and calls a decision rule ultrapessimistic about the experiment X if X is "relevant," yet has no value. An experiment is called relevant if there exist distinct states ω_1 and ω_2 , and a set $S \in \mathcal{S}$, such that

$$\mathbb{P}(X \in S | \omega_1) \neq \mathbb{P}(X \in S | \omega_2).$$

As a corollary, Parmigiani shows that as long as there are at least three states, and nature has an equilibrium strategy that is supported on a strict subset of the state space, then there exist experiments for which the minimax rule is ultrapessimistic. The need for three distinct states stems from a feature of games in regret form: so long as the decision maker does not have a dominant strategy, nature must employ a mixed strategy in equilibrium. Thus, nature's equilibrium strategy is supported on at least two states, and ultrapessimism cannot appear in the regret form of a decision problem with two states.

This hints at a more significant distinction between minimax negative income and minimax regret. In the negative income game, nature may have strictly dominated strategies, and information in favor of these states will always be worthless. In a regret form game, no state ever represents a strictly dominated strategy for nature. This ensures that valuable experiments always exist, and further, a strong enough signal in favor of a given state will decrease the regret obtained *in that state*.

Proposition 2. *Consider a decision problem with finitely many states and finitely many primary acts. If the minimax regret is greater than zero, there exists an experiment with strictly positive value to a minimax regret agent.*

Note that the experiment need not reveal the true state with certainty. Further, Savage's (1954) analysis of partition problems demonstrates that repeated observation will cause the minimax regret to approach zero.

Proof. Let \bar{L} denote the maximum regret from any act in any state, and let L^* denote the minimax regret of the decision problem in the absence of an experiment. Label the states of nature $\omega_1, \omega_2, \dots, \omega_n$, and choose the experiment X such that

$$\mathbb{P}(X = i | \omega_i) = \frac{\bar{L} - L^*}{\bar{L}}, \quad \mathbb{P}(X = i | \omega_j, j \neq i) = \frac{L^*}{(n-1)\bar{L}}.$$

It is a simple exercise to verify that this experiment enables a strict improvement in minimax regret, and the signal i must increase the probability that the optimal action in state i is chosen. \square

I contend that the spirit of Savage’s objection to minimax negative income lies in the existence of decision problems for which no such experiment can be devised, and in this sense, the critique clearly does not apply to the minimax regret decision rule. Indeed this claim is bolstered by the observation that a Bayesian decision maker is vulnerable to the charge of ultrapessimism as defined above. Recall the payoff table from Savage’s example.

	B_1	B_2
f_1	−1	−1
f_2	−10	1

Given the same experiment X described in section 2, a Bayesian with a prior that assigns equal probability to B_1 and B_2 will place zero value on the experiment so long as $g \geq \frac{2}{11}$. Clearly, experiments with $\frac{2}{11} \leq g < \frac{1}{2}$ are relevant according to the definition given by Parmigiani (1992), so this would have us label the Bayesian as ultrapessimistic. Further, recall from our previous analysis that for this range of g , a minimax regret agent attaches a strictly positive value to the experiment. The existence of “relevant” but worthless experiments for a minimax regret agent fails to distinguish minimax regret from other decision rules. The existence of decision problems in which no experiment is valuable stands out as a property of minimax applied to negative income.

4 Discussion

The analysis from the previous section suggests that the relevance of an experiment to a decision maker is best viewed in the context of the decision rule being applied: we could simply say that an observation is relevant if it has value to the decision maker, given the decision rule. The real question becomes what observations are relevant. For any pair of decision rules we have considered, we can construct examples in which an observation is relevant to one but not the other, or vice versa.

Minimax applied to negative income has the peculiar feature of ignoring any conceivable experiment in some decision problems. Other examples of this are presented by Hodges and Lehmann (1950) and by Radner and Marschak (1954). Hodges and Lehmann examine the problem of estimating the parameter of a binomial distribution. Under one loss function they consider, the minimax negative income estimate is constant despite any observations made. Similarly, Radner and Marschak posit a hypothetical gamble on the outcome of flipping a weighted coin. They show that a decision maker applying a minimax negative income rule never observes more than one flip of the coin to obtain information, even if doing so is nearly costless. More recently, Manski (2004) has noted this problem in the context of treatment choice. In all of these examples, the minimax regret rule avoids the pitfall of failing to utilize

information; Savage (1954, p. 200) even addresses the example of Hodges and Lehmann (1950) directly.

The Hurwicz criterion has been proposed as a way to address the pessimism of minimax applied to negative income. In its general form, the criterion suggests choosing an action f to maximize

$$H(f) = \phi(\sup_{B_i} I(f; B_i), \inf_{B_i} I(f; B_i)),$$

where ϕ is any monotonically increasing function of both arguments, and $\{B_i\}$ is a partition of the state space. Unfortunately, as shown by Radner and Marschak (1954), the Hurwicz criterion falls victim to essentially the same problem as the minimax negative income rule. A small set of states can entirely dominate the decision making process, and there are examples in which no conceivable experiment is relevant. In the simplest case, suppose there are three possible events, B_1 , B_2 , and B_3 . Suppose that if B_1 obtains, then no matter what action is chosen, the payoff will be higher than in all other events, and suppose that if B_3 obtains, then no matter what action is chosen, the payoff will be lower than in all other events. The Hurwicz criterion then simplifies to choosing an action which maximizes

$$H(f) = \phi(I(f; B_1), I(f; B_3)),$$

and the consequences of an action in event B_2 are completely ignored, even if available information strongly suggests that B_2 obtains.

As mentioned in the introduction, treatment choice has recently become an important domain for the application of minimax estimation procedures.¹ Early work in this area by Manski (2004, Sec. 2.3) observed the problem of ultrapessimism when employing a minimax negative income rule. In the illustrative example he provides, the issue is precisely that nature has a dominant strategy—namely, choosing an especially poor outcome distribution for the innovative treatment—resulting in a no-data rule. This demonstrates that the problem with the minimax negative income rule extends beyond hypothetical examples to contexts of serious practical importance. Minimax regret is free from this flaw, allowing meaningful application of minimax procedures to treatment choice.

In this same context, Stoye (2009, 2012) has renewed concerns about the pessimism of minimax applied to regret while simultaneously offering solutions. In particular, without any restriction on the space of strategies available to nature, information provided by covariates is of no use to a minimax regret procedure: the procedure will treat the outcome distribution for each value of the covariate independently, as though it were completely unrelated to the distributions for other covariate values. Intuitively this is clear: in the zero-sum game, nature refrains from providing extra information to the statistician by choosing the distributions for each value of the covariate independently from one another. Stoye (2009) has noted that with infinitely many covariate values, this leads to a no-data rule for the minimax regret procedure. This result should be unsurprising. Given the state space under consideration, nature is free to independently choose a different outcome distribution for infinitely many values of the covariate. Under these circumstances, a Bayesian estimator is inconsistent

¹Manski (2004, 2007); Stoye (2009, 2012)

essentially everywhere². Finite samples are unable to provide meaningful information in these situations. Stoye (2012) takes the view that this problem results from an underspecification of the prior information, and he suggests incorporating prior information into the model by restricting the space of strategies available to nature.

Savage (1954, p. 203) acknowledged a distinct but related criticism of minimax decision rules, derived from another example given by Hodges and Lehmann (1950). The example again concerns estimating the parameter of a binomial distribution. If squared error is used as the loss function, then the minimax regret estimate suffers higher loss than the maximum likelihood estimate over most of the parameter space. As the number of trials grows without bound, the region in which minimax regret is superior becomes vanishingly small. In this example, minimax appears overly concerned with the possibility that the parameter lies in a small neighborhood of $\frac{1}{2}$, and it sacrifices performance elsewhere to reduce the risk of loss in this part of the parameter space. This phenomenon could also be described as pessimism. Stoye (2009) notes what I believe is a related example in his concluding remarks, in which a sample that is heavily skewed towards one treatment leads to a counter-intuitive recommendation.

Such pessimism may be an inevitable consequence of ensuring “uniform” performance across all possible states. In the case of the binomial distribution, it is a result of the relatively high variance of a binomial random variable with parameter close to $\frac{1}{2}$. This feature of the binomial distribution renders it more difficult to estimate the parameter in this region, and obtaining uniform performance necessitates sacrificing performance in regions with lower variance. One might consider rescaling the loss function to obtain a more desirable estimate. Indeed, the choice of loss function for parameter estimation seems to have received little attention, yet it is fundamental to the behavior of the minimax regret procedure. With modern computational tools, it may be feasible to consider minimax regret estimates over a range of loss functions, providing another potential avenue for future research.

References

Freedman, D. (1965), “On the Asymptotic Behavior of Bayes Estimates in the Discrete Case II.” *The Annals of Mathematical Statistics*, 36, 454–456.

²Consistency in this context means that the estimator will converge in probability to the true parameter value. A pair comprised of a prior distribution on the parameter space, and a true value of the parameter, is consistent if that particular prior will converge in probability to that particular parameter value as the amount of available data increases. Freedman (1965) gives an example with countably many parameters where the set of consistent prior-parameter pairs is of category 1. More generally, as a consequence of well-known results in measure theory (See Parasarathy, 1967, p. 29) and topology (See Schaefer, 1966, p. 23), all measures on complete, separable metric spaces are tight, and hence put probability one on a countable union of compact sets. Moreover, all locally compact Hausdorff topological vector spaces are finite-dimensional. This implies that all measures on an infinite-dimensional Hausdorff topological vector space assign probability one to a category 1 (“meagre”) set. Therefore, with infinitely many parameters, a Bayesian estimator is inconsistent outside an insignificant portion of the parameter space.

- Hodges, J. and E. Lehmann (1950), “Some Problems in Minimax Point Estimation.” *Annals of Mathematical Statistics*, 21, 182–197.
- Manski, C. (2004), “Statistical Treatment Rules for Heterogeneous Populations.” *Econometrica*, 72, 1221–1246.
- Manski, C. (2007), “Minimax-Regret Treatment Choice with Missing Outcome Data.” *Journal of Econometrics*, 139, 105–115.
- Parasarathy, K. (1967), *Probability Measures on Metric Spaces*. Academic Press.
- Parmigiani, G. (1992), “Minimax, Information, and Ultrapessimism.” *Theory and Decision*, 33, 241–252.
- Radner, Roy and Jacob Marschak (1954), “Note on Some Proposed Decision Criteria.” In *Decision Processes* (R. Thrall, C. Coombs, and R. Davis, eds.), 61–69, John Wiley and Sons.
- Savage, L. (1954), *The Foundations of Statistics*, second edition. Dover Publications Inc. 2nd Ed. published 1972.
- Schaefer, H. (1966), *Topological Vector Spaces*. MacMillan.
- Stoye, J. (2009), “Minimax Regret Treatment Choice with Finite Samples.” *Journal of Econometrics*, 151, 70–81.
- Stoye, J. (2011), “Statistical Decisions under Ambiguity.” *Theory and Decision*, 70, 129–148.
- Stoye, J. (2012), “Minimax Regret Treatment Choice with Covariates or with Limited Validity of Experiments.” *Journal of Econometrics*, 166, 138–156.
- Wald, A. (1950), *Statistical Decision Functions*. John Wiley and Sons.