PROFILING PROBLEMS WITH PARTIALLY IDENTIFIED STRUCTURE

by

William A. Brock[1]
Department of Economics
University of Wisconsin, Madison
First Draft, October 23, 2003, This Draft, November 30, 2005

ABSTRACT

This paper presents two classes of profiling problems where the relevant structural inputs to the problems may be only partially identified: (i) Class I is where profilees are not "strategic" in altering their characteristics; (ii) Class II is where profilees strategically alter their characteristics at a cost. This paper studies Class I profiling problems. A very brief study of Class II profiling problems is conducted. The analysis can be unified by assuming there is a utilitarian planner that attempts to optimize the sum of welfare. We compare maximin solutions and minimax regret solutions in two contexts. The first context we call "bottoms up" where minimal assumptions are made on unknown structures. The second context we call "top down" because stronger assumptions based upon economic theory or some such device are made. Connections are made to the literature on zero sum game theory. An unexpected feature of the minimax regret solution is that a minimax regret planner may end up learning unknown features of the population it faces rather quickly.

1. INTRODUCTION

This paper presents two classes of profiling problems where the relevant structural inputs to the problems may be only partially identified: (i) Class I is where profilees are not "strategic" in altering their characteristics; (ii) Class II is where profilees strategically alter their characteristics at a cost. This paper studies Class I profiling problems. A very brief study of Class II profiling problems is conducted. The analysis can be unified by assuming there is a utilitarian planner that attempts to optimize the sum of welfare over treatments.

We compare maximin solutions and minimax regret solutions in two contexts. The first context we call "bottoms up" where minimal assumptions are made on unknown structures. The second context we call "top down" because stronger assumptions based upon economic theory or some such device are made. Connections are made to the literature on zero sum game theory.

The maximin solution maximizes welfare over treatments assuming the worst case bound on partially identified structure. The minimax regret solution minimizes regret over treatments. Maximum regret at a particular treatment is the maximum difference over possible structures between the optimal value of treatment for a particular structure and the actual value of the particular treatment at that structure. The minimax regret solution takes into account the best case bound on partially identified structure and, in many scalar cases, randomizes with a probability equal to the ratio of the best case to the spread of the best case and the worst case. This property of the minimax regret solution is a nice property that helps deal with criticisms of maximin for being too conservative. An unexpected feature of the minimax regret solution is that a minimax regret planner may end up learning unknown features of the population it faces rather quickly (Brock (2004)).

In section 2, we lay out a basic framework for Class I problems and develop a leading example of a lender who profiles across observable type strata. The lender's problem is used as an expository vehicle. In section 3, we draw connections between the maximin and minimax regret problems and zero sum game theory. This connection suggests tools of analysis as well as potential extensions by exploiting the large literature on zero sum games, and adaptive learning in such contexts.

Section 4 conducts what we call a "top down" approach in contrast to the earlier sections which conducted what we call a "bottoms up" approach. In the bottoms up approach, the set of possible alternative structures is large. We call that set the "ambiguity set." The "top down" approach replaces the "large" ambiguity set with "small" ambiguity sets centered around a "baseline" This baseline has some claim to credibility based upon, perhaps, previous empirical work as well as theory. This approach enables us to draw a connection between the rapidly growing area of research on robustness and our concerns here. We find that there are consistent patterns in the directions where the planner needs to robustify treatment choice more intensely as a function of the structure of the ambiguity set of partially identified structures.

Section 5 briefly applies the tools developed in previous sections to Class II profiling problems where profilees have a strong incentive to falsely report their type to the planner. We restrict ourselves to a subclass of these problems where the cost function to false reporting is "L-shaped" in the sense that a profilee can report a type up to "e" greater than its true type for a cost of zero, but the cost becomes infinite if it reports a type that deviates above its true type by more than "e". We sketch how the tools can be applied to this class of problems and sketch what the maximin and minimax regret solutions look like.

In section 6, we present a brief sketch of a problem of detecting discrimination by employers when employers have to conduct decision making under ambiguity about unfamiliar types. This example is included not only because it is of current policy importance but also to suggest potential applications of the tools developed in earlier sections and to show the similarity of analysis across a potentially wide set of applications.

Finally, Section 7 contains a summary, conclusions and suggestions for future research.

2. BASIC FRAMEWORK

The basic framework of the paper builds on Manski's articles, (2004a,b), on treatment rules, his book (2003) on partially identified structures, and Brock and Durlauf (2004) on local robustification. We also suggest applications of ideas on dealing with model uncertainty from Brock, Durlauf, and West (2003), (2004). We use Manski's (2004a) notation as much as possible. There is a planner who wishes to choose a policy t in T (i.e. a "treatment rule") to maximize population mean welfare. Manski assumes that the welfare from assigning treatment t to person j is $u[y_j(t),t,x_j]$ where $u(.,.,.):YxTxX\text{-->}R$ is the welfare function which is "individualistic", i.e. u does not depend upon the outcomes of persons other than j and x_j denotes observable characteristics of person j. Hence the problem can be analyzed term by term, i.e. for each covariate x.

Fix a particular covariate x in X and suppress it in the notation for now. We change Manski's (2004a,b) notation so that it is compatible with the general framework to be described below. Let T be the treatment set, $T=\{1,2,...,n\}$, let $\pi=(\pi(1),...,\pi(n))$ denote the vector of $\pi(t):=E\{u(y(t),t)\}$.

Manski (2004b) treats a case where,

$$(2.1)\quad \pi(t)=E\{u(y(t),t)|z_t=0)\}P(z_t=0)+E\{u(y(t),t)|z_t \text{ not } 0)\}P(z_t \text{ not } 0),$$

where "$z_t=0$" ("z_t not 0") stands for "outcome treatment t observed" ("outcome treatment t not observed") and the rest of the notation is self explanatory. The problem is that the scientist can not observe $E\{u(y(t),t)|z_t \text{ not } 0)\}$ so we assume it takes values in the closed interval $[a(t),b(t)]$, which we call the "ambiguity set.". Thus $\pi(t)$ can be any value in the closed interval $[\pi_*(t),\pi^*(t)]$ where the lower and upper limits are defined by (2.1) and $[a(t),b(t)]$.

By borrowing from Manski (2004a,b) we now have enough machinery to formulate the minimax regret planner's problem, the maximin planner's problem, and formulate a pedagogical example of a lender who is lending to a partially identified population. The lender's problem will be used to illustrate differences between the minimax regret and maximin objectives. We will find that, in many cases, minimax regret planners learn partially identified features of their populations much faster because they may find it optimal to conduct a form of experimentation which relates to mixed strategies in game theory.

With our notation now in place, the maximin problem and minimax regret problem for each covariate (we suppress x in the notation) can be stated thus,

$$(2.2a)\quad \max_t \min_{\pi(t)}\{\pi(t)\}=\max_t\{\pi_*(t)\},$$

$$(2.2b)\quad \min_d \max_\pi\{\max_u \Sigma u(t)\pi(t)-\Sigma d(t)\pi(t)\},$$ where Σ is over t, where u,d range over the n-dimensional simplex S_n.

As in Manski (2004b) note that the $\max_u$ takes place at extreme points of the simplex S_n, thus, the simplex S_n can be written as a union over t of subsets S_{tn} where t is optimal. On these subsets, S_{tn} we observe that the largest value of

$$(2.3)\quad (1-d(t))\pi(t)-\Sigma' d(s)\pi(s) \text{ is } (1-d(t))\pi^*(t)-\Sigma' d(s)\pi_*(s),$$

where Σ' denotes summation over all s not equal to t. Thus minimization of maximum regret collapses to the problem,

$$(2.4)\quad \min_d \max_t\{(1-d(t))\pi^*(t)-\Sigma' d(s)\pi_*(s)\},$$

which is the minimum of an envelope described by a family of n linear functions of the vector d in S_n. The envelope is convex and this reduces the complexity of the problem. The same argument applies for each value of the covariate x in X.

A special case is T={1,2}. Since d(1)+d(2)=1, (2.4) boils down to

$$(2.5)\quad \min_d \max\{(1-d(1))[\pi^*(1)-\pi_*(2)],d(1)[\pi^*(2)-\pi_*(1)]\}.$$

One computes the minimax regret solution by listing the possibilities for the minimum of two lines in (2.5), i.e. one just lists the possibilities for an envelope generated by the two lines each of which can have positive or non-positive slope.

If we have three treatments, the two one dimensional lines in (2.12) get replaced by three two dimensional planes and the slopes of the lines get replaced by the direction numbers of the planes. While it is possible to list the possibilities for these direction numbers and, hence, list the possibilities for the envelope, we content ourselves with pointing out that this geometry shows why one should expect "mixed strategies" to be quite common in minimax regret solutions. Karlin (1959) provides techniques for listing out the possibilities and explicitly constructing mixed and pure strategies in closely related structures.

At the risk of repeating, if we now insert x in X into the above notation and results we have Manski's (2004b) result for the minimax regret problem. Notice Manski's emphasis on the possibility of a mixed treatment strategy. Mixing will play an important role in adaptive learning and in the other formulations of dealing with different forms of the planner's ignorance below.

THE LENDER'S PROBLEM

Let us apply this work to a lender's problem. Suppose a lender faces, for each covariate x, a partially identified population of borrowers. Let the "treatments" be denoted 1="do not lend", 2="lend". Suppress covariate x in the notation. Let $\pi(2)$ denote unknown net return which is assumed to lie in the set $[\pi_*(2),\pi^*(2)]$. Let $\pi(1)=0$ because return is zero if t=1, i.e. "do not lend" is chosen, and let t, d, π denote vectors of "treatments" in T={1,2} and payoffs. Then the maximin problem is given by

(2.6) $\max_d \min_\pi (d\pi) = \max_d \min_\pi \{d(1)\pi(1)+d(2)\pi(2)\} = \max_d \{d(2)\pi_*(2)\}$.

Hence the RHS is zero with d(2)=0, if $\pi_*(2)\le 0$ and the RHS is $\pi_*(2)$ with d(2)=1, if $\pi_*(2)>0$. Hence, bringing the covariate x back into the notation, if we assume x is scalar, $\pi_*(2,x)$ is strictly increasing in x and continuous in x, then the maximin lender conducts a simple profiling strategy: Lend to $x>x^*$ where x^* is determined by the equation $\pi_*(2,x^*)=0$. This is a very conservative worst case strategy.

Turn now to minimax regret. The minimax regret problem is given by

(2.7) $\min_d\{\max_\pi[\max_t(t\pi)-d\pi]\}$.

But by (2.5) above, we have,

(2.8) $\min_d \max_\pi[\max_t(t\pi)-d\pi] = \min_d \max\{d(2)(-\pi_*(2)),(1-d(2))\pi^*(2)\}$,

since $\pi^*(1)=\pi_*(1)=0$.

Proceeding as in Manski (2004b),it is easy to see that $\pi^*(2)\le 0$ implies d(2)=0 (we take d(2)=0 in the boundary case $\pi^*(2)=0$). Furthermore $\pi_*(2))>0$ implies d(2)=1. We get a mixed strategy for the case $\pi_*(2))<0$ and $\pi^*(2)>0$, i.e.,

(2.9) $0<d(2)=\pi^*(2)/[\pi^*(2)+(-\pi_*(2))]<1$.

All of the above generalizes to lending problems with covariates x in X by writing $\pi_*(2,x)$ and $\pi^*(2,x)$ and repeating the above analysis "x by x".

Let us explain (2.9) further. Notice that the denominator is the "width" of the ambiguity set while the numerator is the best case profit. Thus the larger is the ratio of the best case profit to the width of the ambiguity interval the larger the fraction of x's that get loans. Notice that ex ante, all members of the x stratum have equal probability of getting a loan. But ex post only the "lucky" ones get a loan. This may seem capricious and unfair to the "losers" in the "loan lottery" conducted by the minimax regret lender. Similar issues of "fairness" of minimax regret "mixed" strategies are sure to come up in policy applications. See Manski's paper (2004c) in this Feature for police strategies and see Manski (2004a) as well as Durlauf's paper (2005) in this Feature for a discussion of fairness issues.

Turn now to another feature of minimax regret mixed strategies. In a simple model of adaptive learning, Brock (2004) shows that a minimax regret lender will learn the unknown profit potential of x's where $\pi_*(x,2)<0$, and $\pi^*(x,2)>0$ in just one lending period, in contrast to the maximin lender. In the case of lending this may be considered a desirable feature of minimax regret lenders in contrast to maximin lenders. However, note that maximin lenders treat all types "fairly". I.e. either all x in a stratum get a loan or none do from a maximin lender. No "lotteries" are ever conducted by a maximin lender. See Brock (2004) for the details.

In the above formulation of a lender's problem the only unknown feature of the population for each x is the unknown payoff to those x's whose outcomes are unobserved. The above formulation attempts to formalize different levels of past experience (in an analytically tractible framework where a notion of "profiling on x" appears) on the part of the lender with lending to different x's.

Turn now to problems where other features of the population are unknown. We shall also expose a close connection between general treatment effect problems and classical zero sum game theory. Since we are using the treatment choice framework to model profiling problems, therefore there is a close connection between profiling problems and classical zero sum game theory.

3. PLANNER'S PROBLEM WITH UNKNOWN TYPE DISTRIBUTIONS

In this section we develop a close connection between treatment choice problems and classical zero sum matrix game theory. This connection seems to us to be under appreciated in the literature. Fix covariate x in X and suppress it in the notation for the moment. But let there be unobserved types k=1,2,...,K and consider the general treatment choice problem, (Manski (2004a, p. 1225)),

(3.1) $\text{Maximize}_t\{\sum_k E(u(y(t),t,k)|k)p(k)\}$, $\sum$ is over k=1,2,...,K types,

where the types probabilities are not known but the payoffs

(3.2) $a(t,k):=E(u(y(t),t|k)$

are known. This is a finite mixture problem that is somewhat different from the finite mixture problems that appear in the treatment choice literature. But it is close enough to that literature for it to give us analytical guidance.

If a decision maker faces a problem of this form, how should it play? If it plays maximin, this is a zero sum matrix game problem of the form analyzed by Karlin (1959), i.e. let the primary player, the maximizing player, be denoted Player I and the minimizing player be denoted Player II. Then Karlin shows that the |T|xK finite matrix game

(3.3) $\mathrm{Max}_u \mathrm{Min}_p \{u'Ap\}$, where A is the matrix of a(t,k)'s,

u is in the compact convex |T| dimensional simplex, p is in the K dimensional compact convex simplex, has a value v. I.e. it has a saddle point where the maxmin=minmax theorem of zero sum game theory holds. The presence of equilibrium mixed treatment strategies u* is general. Karlin gives explicit solution techniques for 2xK and |T|x2 games which are still somewhat tractible for 3xK and |T|x3 games. He also gives solution techniques for |T|xK games where n=|t|=K. We develop this below as well as some theory for minimax regret solutions.

PLANNER FACES AN UNKNOWN, BUT TREATMENT INDEPENDENT, TYPE DISTRIBUTION

Consider a lender who is facing two types of borrowers. It knows the net payoff to each type of borrower but it does not know the probabilities of each type in the population it faces. Here t=1 denotes "do not lend" and t=2 denotes "lend". There is a processing cost c(2) if lending is done. Type 1 has zero gross profit potential, i.e., A(2,1)=0, and type 2 has positive net profit potential, i.e. A(2,2)-c(2)>0. The lender does not know p(1) and p(2), but it does know {A(t,j), c(t)}. We show below how the lender might proceed.

Recall the problem (3.3) for the special case |T|=2,

(3.4) $\mathrm{Max}_u \mathrm{Min}_p \{u'Ap\} = \mathrm{Max}_u \mathrm{Min}_j \{L_j(u)\}$ where,

following Karlin (1959, pp. 41-42), the line $L_j(u)$ is given by

(3.5) $L_j(u) = u(1)A(1,j) + (1-u(1))A(2,j)$,

the lower envelope of the set of K lines $\{L_j(u)\}$ is concave (cf. Karlin (1959, Figure 2.1)), and the maximizer u maximizes the concave function, $\mathrm{Min}_j\{L_j(u)\}$. Much as in the above analysis of minimax regret, we can divide the analysis into Case (i): The maximum is at the boundary of [0,1], Case (ii) The maximum is in the interior of [0,1]. The interior maximum is a mixed strategy. The reader may now follow Karlin (1959) to complete the analysis and extend it to cases of 2xK and 3xK games and |T|x2 and |T|x3 games. We wish to concentrate on minimax regret solutions.

The minimax regret problem is given by

(3.6) $\mathrm{Min}_d \mathrm{Max}_p \{\mathrm{Max}_u u'Ap - d'Ap\}$

$$= \mathrm{Min}_d \mathrm{Max}_p \mathrm{Max}_k \{(1-d(k))\textstyle\sum_j A(k,j)p(j) - \sum_i{}' d(i)(\sum A(i,j)p(j)\}:=\mathrm{RHS},$$

where $\sum_i{}'$ is the sum over all i not equal to k. The last equality follows from recognizing that maximization over u in the |T| dimensional simplex takes place at the extreme points of that simplex.

We gain insight into the structure of the problem by considering the special case |T|=2, K=2. In this case, the expression in {.} becomes, for k=1,

(3.7) $(1-d(1))[A(1,1)p(1)+A(1,2)(1-p(1))]-d(2)[A(2,1)p(1)+A(2,2)(1-p(1)]$

$=(1-d(1))[Dp(1)+b]$, where

(3.8) $D:=A(1,1)-A(1,2)-(A(2,1)-A(2,2))$, $b:=A(1,2)-A(2,2)$.

Similarly, using again that $1=d(1)+d(2)$ and $1=p(1)+p(2)$, we may write the expression in {.} in (3.6) for k=2 as,

(3.9) $(1-d(2))[A(2,1)p(1)+A(2,2)(1-p(1)]-d(1)[A(1,1)p(1)+A(1,2)(1-p(1))]$

$=d(1)[-Dp(1)-b]$.

Hence, RHS (3.6) becomes,

(3.10) $RHS=Min_d Max_p Max\{(1-d(1))[Dp(1)+b],d(1)[-Dp(1)-b]\}$

Put $\pi:=Dp(1)+b$, and rewrite (3.10) as

(3.11) $Min_d Max_\pi Max\{(1-d(1))\pi,-d(1)\pi\}$.

The inner maximization over π takes place at the extreme points of the set $\Pi=[\pi_*,\pi^*]$, so we now choose d(1) to solve the problem,

(3.12) $Min_d Max\{(1-d(1))\pi^*,-d(1)\pi_*\}$.

For the "main" case, $\pi_*<0,\pi^*>0$, this gives the solution, which we have seen before,

(3.13) $d(1)=\pi^*/(\pi^*+(-\pi_*))$.

With this background work, return to the problem of the lender who does not know p(1) and p(2), but who does know {A(t,j), c(t)}. What should it do? If it does maximin, it solves

(3.14) $max_u min_p\{u'Ap-u'c\}=max_u min_p\{u(2)[A(2,2)p(2)-c(2)]\}$,

which gives the solution $u(2)=0$ for every $c(2)>0$, no matter how high A(2,2) is. I.e. no matter how high the profit potential is to lending to type 2's, and no matter how low the processing cost per loan, c(2) is, (so long at it is positive) the maximin lender assumes the worst case scenario, i.e., $p(2)=0$.

This is a version of the recurring theme of extreme pessimism of maximin. Of course the extreme pessimism of maximin is "toned down" in the presence of prior information such as knowledge of a lower bound on p(2) that is positive. Turn now to minimax regret.

The minimax regret problem is given by (omitting familiar steps),

(3.15) $min_d max_p max\{(1-d(2))[A(2,2)p(2)-c(2)],-d(2)[A(2,2)p(2)-c(2)]\}$.

Application of the above solution technique gives

(3.16) $d(2)=[A(2,2)-c(2)]/A(2,2)$,

which is a "net return" measure on treatment 2 measured relative to "gross return" $A(2,2)$. This seems more sensible than the maximin solution of zero for any $c(2)>0$, no matter how small, and any $A(2,2)>0$, no matter how large. This point is related to Manski's (2004a) point that maximin "ignores" the data whereas here the point is that maximin "ignores the return". The minimax regret lender "samples" the population in the fraction $d(2)$ which increases as type 2 net profit, $A(2,2)-c(2)$, increases relative to type 2 gross profit, $A(2,2)$.

This minimax regret lender will learn the relevant features of the population it faces within one period of results from lending, whereas the maximin lender will not learn at all (Brock (2004)). In the language of profiling the maximin lender profiled out all borrowers in this population, whereas the minimax regret lender randomly picked people and gave them loans with probability $d(2)$ in (3.16).

Notice that a "fairness" issue is raised here because equals were not treated equally ex post by the minimax regret lender even though they were treated equally ex ante. I.e. each borrower had an equal probability of winning the "lending lottery" conducted by the minimax regret lender. Turn now to a "top down" approach to dealing with optimal action under partially identified structural characteristics.

4. A "TOP DOWN" APPROACH

LOCAL ROBUSTIFICATION OF "IDEALIZED" TREATMENT RESPONSE

The lender's problem is a useful introduction to the more general problem we wish to address now.

Consider Manski's (2004a, Section 2.1) "idealized" treatment response problem, where there is a planner whose goal is to maximize population mean welfare which is achieved by

(4.1) $U^*(P):=\sum p(x)\{\max\{Eu[y(t),t,x]|x\}\}:=\sum p(x)\{\max\{\pi(t,x)\}\}$,

where $\sum$ is over x in X, $\pi(t,x):=E\{u[y(t),t,x]|x\}$, max is over t in $T:=\{1,2\}$, and I change the notation slightly from Manski's to ease typing. It is clear that the lender's problem is a special case of this framework. Suppose now that the planner faces a new problem whose structure is "close" to the above problem in some metric. How should the planner behave? We develop a strategy for the planner below. Put, for the "baseline problem" above,

(4.2) $a_o(x_i):=E\{u[y(2),2,x_i]|x_i\}-E\{u[y(1),1,x_i]|x_i\}, i=1,2,\ldots,n,$

for $X:=\{x_1,x_2,\ldots,x_n\}$.

It is easy to see that the optimum for the baseline problem is to choose treatment $t=2$ when $a_o(x_i)>0$ and choose treatment $t=1$ when $a_o(x_i)<0$. Assume $a_o(x)$ is increasing and continuous in x so there is a unique x^* such that $a_o(x^*)=0$ and all x's$>x^*$ are assigned to treatment 2. Profiling appears here because the planner profiled on x and chose the optimal cutoff x^* for $t=2$ over $t=1$. At the risk of repetition for clarity, the cutoff, x^*, was determined by the rule

(4.3) $a(x):=E\{u[y(2),2,x]|x\}-E\{u[y(1),1,x]|x\}>0$, for $x>x^*$.

Hence, the planner's objective (e.g. the lender's objective in this very special case of (4.1)), can be written as

(4.4) $J(u,a):=\sum u_i a(x_i)p(x_i)$.

Here $u_i=1$ if the planner gives treatment 2 to x_i and zero otherwise.

In many binary treatment applications it is natural to take t=1 as "status quo" and (4.4) as the planner's objective. Suppose the planner knows that the population it faces is structurally near the baseline population where objective function (4.4) is appropriate. Suppose the planner has no reason to suspect that the new population structurally deviates in any particular direction. In this case the planner can locate the directions of structural deviation against which it is most vulnerable (and find optimal strategies against those directions) by introducing a fictional Adversarial Agent and computing Nash equilibria for an appropriate non cooperative game. I.e. it finds the Nash equilibria, u^*,a^*, defined by the following conditions, where the planner is the maximizing player and the Adversarial Agent is the minimizing player,

(4.5) $J(u^*,a^*)\geq J(u,a^*)$, $u:=(u(x_1),...,u(x_n))$, $u(x_i)$ in $\{0,1\}$,

(4.6) $J(u^*,a)\leq J(u^*,a^*)$, for all a such that $||a-a_o||\leq e$,

where $e>0$ is appropriately small and $||.||$ is an appropriate distance measure between $a:=(a(x_1),...,a(x_n))$ and $a_o:=(a_o(x_1),...,a_o(x_n))$. It is important to note that this non cooperative game and the fictional Adversarial Agent are just expository devices to obtain the planner's best strategy for protecting iself against mis-specifications of the structure of the objective and constraints that it faces. We call the procedure of formulating this game and solving for its Nash equilibria, a "local robustification analysis."

LOCAL ROBUSTIFICATION FOR "IDEAL" TREATMENT RESPONSE PROBLEMS

We have now developed enough machinery that it is appropriate to briefly sketch what a general local robustification analysis of (4.1) would look like. Recall (4.1) and rewrite it as follows,

(4.7) $U^*(P):=\sum p(x)\{\max\{Eu[y(t),t,x]|x\}\}:=\sum p(x)\{\max\{\pi(t,x)\}\}$

$=\sum p(x)\{\max\{\sum d(t,x)\pi(t,x)\}\}\}$

$=\sum p(x)\{\pi(1,x)+\max\{\sum' d(t,x)[\pi(t,x)-\pi(1,x)]\}\}\}$

$:=\sum p(x)\{\pi(1,x)+\max\{\sum' d(t,x)[a(t,x)]\}\}\}$,

where "max" is over $\{d(1,x),d(2,x),...,d(K,x)\}$ such that $\sum d(t,x)=1$, $d(t,x)\geq 0$, for all t in $T=\{1,2,...,K\}$, for all x in X, and $\sum'$ is over $t=2,...,K$. It is easy to see that maximizing over K choices is the same as maximizing the above problem over the K simplex in the nonnegative orthant of K dimensional Euclidean space. This formulation is convenient for the analysis that follows which exploits the J(u,a) framework used earlier.

We take the position here that the main interest in robustification is to robustify the treatment decision d(t,x) against misspecification of the structure of the problem. We see from (4.7) that the structure that matters from this position is the collection of functions {a(t,x), t in T, x in X}. Hence we map this problem into the J(u,a) framework by putting u=d.

There is the analytical issue of which norm, $||.||$, one should choose. The Euclidean norm is harder to work with, so we choose the maximum norm here. A first draft of this paper (Brock (2004)) contains results using the Euclidean norm and draws connections between the structure of Euclidean norm robustification results and hazard functions in duration analyses.

Consider, the objective function in (4.4),

(4.8) $J(u,a):=\sum u_i a(x_i)p(x_i)$.

The maximin problem, and its solution, using the maximum norm to define an e-neighborhood, $N:=\{a|\max_i\{|a(x_i)-a_o(x_i)|\le e\}\}$ is given by,

(4.9) $\max_u \min_a \sum u_i a(x_i)p(x_i)=\max_u \sum u_i (a_o(x_i)-e)p(x_i)$,

(4.10) $u^*(x_i)=1[a_o(x_i)-e\ge 0]$.

Notice the additive separability that is "preserved" using the maximum norm to define e-neighborhoods in contrast to using the Euclidean norm to define e-neighborhoods (Brock (2004)).

This useful additive separability of maximum norm e-neighborhoods allows us to do the local minimax regret problem "x by x". Thus we can directly apply our above analysis of minimax regret to conlude that planner should follow the rule: (i) Choose $u^*(x)=0$ if $a_o(x)+e\le 0$; (ii) Choose $u^*(x)=1$ if $a_o(x)-e\ge 0$; (iii) Choose u=1 with probability $d^*(x)$ given by

(4.11) $d^*(x)=[a_o(x)+e]/[a_o(x)+e-(a_o(x)-e)]=[a_o(x)+e]/(2e)$,

when, $a_o(x)+e>0$ and $a_o(x)-e<0$. Notice that $d^*(x)=1/2$ when $a_o(x)=0$. Also notice that, RHS (4.11) is a probability because the size of $a_o(x)$ is constrained by the relations $a_o(x)+e>0$ and $a_o(x)-e<0$. Indeed, the relations $a_o(x)+e>0$ and $a_o(x)-e<0$, can be used to show that RHS (4.11) is bounded between zero and 1.

We have done enough to show that it should be possible to replace any local maximin approach to local robustification by a local minimax regret approach to local robustification. Using the minimax regret objective rather than the usual maximin objective in robustification analysis seems desirable because it takes into account the best case scenario as well as the worst case scenario and mixes according to the ratio of the best case payoff to the width of the ambiguity interval in payoffs. This seems more sensible than simple robustification against the worst case payoff as in maximin robustification. We also have indicated that a strategic choice of the robustification neighborhood can simplify the analysis. It is beyond the scope of this article to do more here. See Brock (2004) for more examples and details. Turn now to classes of problems where the profilees can alter their characteristics at a cost.

5. CLASS II PROFILING PROBLEMS

In class II profiling problems, profilees can alter their characteristic x to x' at a cost C(x,x'). If we assume the distribution of types and C(x,x') is common knowledge to both lenders and lendees, we can study equilibrium concepts. There is a sub class of these problems where our treatment may be brief by mapping the analysis into a special case of the above.

In order to get started let us treat a simplest example first. Suppose $C(x,x')=0$ for $x'\le x+e$, and is infinite for $x'>x+e$, for some $e>0$. Thus an x will always pose as $x'=x+e$ if it is worthwhile to pose at all. Suppose there are only two possible types $x_1<x_2$ with probabilities q_1, $1-q_1$ and this information is common knowledge to both sides. Let $\pi(x)$ denote the net profit of lending to an x and let it be known to the lender. Assume $\pi(x)$ is increasing and continuous in x.

Suppose, at first blush, the lender uses the rule: Lend to an x if $x\ge x^*$ where x^* solves $\pi(x)=0$. Assume $x_1<x^*\le x_2$, so that lender plans to reject x_1, but lend to x_2. If the "bad" type x_1 differs from the "good" type x_2 by no more than e, then x_1 will pose as x_2. But lender knows this and knows that the cost to x_1 of such posing is zero. Hence lender must protect itself by assuming it faces a "pool" of q_1, x_1's and $1-q_1$, x_2's. Mean profits for lender facing this population of lendees is

(5.1) $d(q_1\pi(x_1)+q_2\pi(x_2))$,

where d=1 if lender lends to the group, d=0 otherwise. Thus, lender acts as if it faces a "worst case" scenario where all the bad types pose as good types. I.e. it lends to the "atom" (x_1,x_2) if the expected value of profit to that atom satisfies

(5.2) $q_1\pi(x_1)+q_2\pi(x_2)\ge 0$

Turn now to a more general version of this problem. If the cost function to lendees of posing as $x'>x$ is zero for $x'-x\le e>0$, and is infinite for $x'-x>e$, and it is common knowledge that the covariate distribution, P(x), is a continuous distribution with bounded support, for example, [0,1], and $e<1$, then lender knows that if it announces a cutoff lending strategy x^*, then lendee of type x will pose as type x+a to maximize the chance of getting a loan. Thus there will be an "atom" at x=1 which consists of all types x such that $x+e\ge 1$, because all types x in [1-e,1] can pose as 1's for free. Furthermore lender should expect that all types x in [0,1-e) will pose as $x'=x+e$. Therefore if lender knows $\pi(x)$ and P(x), it should lend to $x'\ge x^*$, for x^* such that,

(5.3) $\pi(x^*-e)\ge 0$

if such x^* exists. If $e\ge 1$, all x's pose as 1's so it lends to the "atom" [0,1] if

(5.4) $\int\pi(x)dP(x)\}\ge 0$.

Hence, we see that if "ambiguity" is present on the size of e in the interval $[0,e^*]$, the lender's problem can be mapped into the general framework developed in this paper. I.e. the maximin lender assumes, for each type x_i' in the set, A, of admissible types, $A:=\{x_1,x_2,\ldots,x_n\}$, it observes, the worst case pool, call it $P_*(x_i')$, that can be formed from the set of admissible types, $A:=\{x_1,x_2,\ldots,x_n\}$ with their probabilities $\{q_1,q_2,\ldots,q_n\}$, and lends provided that the net return on that pool, denote it by $\pi_*(x_i')$, is positive. Notice that if any type x_i "lies" i.e. reports that it is a higher type $x_j>x_i$ in A, that it always "lies to the max" and reports the largest "reachable" type x_j in A that is in the "reachable" set, $[x_i,x_i+e^*]$. Meanwhile the best possible pool $P^*(x_i')$ is just the truth. I.e. the reported x_i' that lender observes is just the true type x_i' and the lender gets the true net return, $\pi^*(x_i')$, on this type. Now we are ready to apply our machinery developed above to this particular subclass of Class II, profiling problems. As we saw from our above work, the maximin lender does not lend to reported type stratum, x_i' if $\pi_*(x_i')<0$. The minimax regret lender, on the other hand, lends to the positive fraction $d(x_i')$ defined by (5.5) below of reported type stratum x_i' provided that, $\pi^*(x_i')>0$,

(5.5) $0<d(x_i')=\pi^*(x_i')/[(-\pi_*(x_i'))+\pi^*(x_i')]$.

Of course this sketch is very brief, but it is enough to show the reader how to set up and analyze a Class II profiling problem using the tools developed in earlier sections.

This formulation can be generalized in many ways.

(i) One can study problems where the cost of posing as $x'>x$ is a convex increasing function, $C(x'-x)$, with $C(0)=0$. The formulation above is enough to suggest that the lender will not want to constrain itself to designing mechanisms where each type is induced to report its true type. Furthermore, suppose the support of x is bounded, W.L.O.G., assume it is the closed interval, $[0,1]$ and this is common knowledge to both sides. We have learned enough from the above work to expect an "atom" to appear at $x=x^*$ where x^* is a cutoff level for lending. This raises issues of existence or nonexistence of "pooling" and "separating" equilibria, analogous to the "screening" literature of the 1970's. Some of these nonexistence problems could be "fixed" by introducing "noisy discrete choice" lenders, producing existence of equilibria indexed by the level of the noise, and taking the noise level to zero to study potential convergence of the equilibria to a "natural" limit. This issue merits study.

(ii) We can study problems where the lender has a wider strategy set. For example let it have access to "atom splitting" strategies where it can randomly audit members of a suspected atom at a cost of S ("S" stands for "search" or "screen") for each. Let the audit detect the true type with probability one if it is conducted and let the punishment cost to the poser if caught be F. Suppose F is fixed by law or custom. If lender could choose F the problem would become trivial because lender should choose an infinite F and audit with infinitesimal probability.

(iii) Potentially more interesting would be to study problems where the possibility of posing interacts with the partial lack of knowledge of $\pi(x)$. For example, the lender could open a trial "loan window" where it accepts N applications from a self reported x-stratum (which may really be x' types) and attempts to learn $\pi(2,x)$ for the population but only gets to observe outcomes for those who accept loans, e.g. those types who report as type x. This is a self selection problem grafted onto a partial knowledge problem. A "structural" econometrician might use economic theory to build a model of this joint process and attempt to use data to estimate that model as in the literature on selection bias, for example. One could, in the standard way, look for sufficient conditions for that model to be identified. However, it may be difficult to get any kind of consensus in the research community that this model is the "right" model. But one might be able to get more consensus on a "cloud" of models surrounding a "baseline" model. If so, then one could adapt the Bayesian Model Averaging approach of Brock, Durlauf, and West (2003).

6. A PROBLEM WHOSE STRUCTURE IS THE SAME AS THE LENDER'S PROBLEM

We used a lender's problem as a main expository vehicle. It might be thought that this is a rather special problem. We briefly develop an employer's problem here to show that the analytic structure is very similar. Furthermore we sketch how our work indicates potentially new insights for the problem of attempting to detect discrimination. This suggests some potentially interesting future research.

For example, suppose there are two types of people, blue and green. An experiment is conducted where fake CV's are constructed with qualifications x for each type where x is varied. These CV's are sent out in response to prospective employers and a response variable is measured, e.g. the number of callbacks received. It is observed that blues receive more callbacks than greens for the same x and it is even observed that the relative callback rate for greens falls as qualifications x increases. What is going on here? Is this evidence of employer discrimination against greens in favor of blues?

Let us use the analytical structure of the lender's problem above to explore this issue. Let treatments be 1 "do not hire", 2 "hire". Let $\pi(2,x,w)$ be the net profit to a firm from a type w employee with characteristics x if hired. The structure of the employer's problem is exactly the same as the lender's problem, except that we add an extra characteristic, "w" to x, where w is blue or green.

We can make the following observations now. First, in order to conclude that the experiment reveals discrimination we need an "invariance assumption". I.e. the hypothetical population constructed by the experimenter must be the same as the actual population faced by the employers. Suppose the experimenter constructed a hypothetical population where the qualifications of blues and greens were exactly the same. If the employer knew this the callback rate should have been exactly the same conditional upon qualifications x. However, invariance is not likely to hold in the field. This lack of invariance is perhaps one reason why the legal notion of "intent" plays such a big role in actual litigation. What can we conclude from our hypothetical experiment?.

We can use such hypothetical experiments plus our work on the lender's problem above to suggest possible alternative reasons for the observed behavior and, possibly design tests to rule out these alternative reasons.

Leading Case: Suppose $\max_x \pi_*(2,x,\text{green})<0<\max_x \pi_*(2,x,\text{blue})$.

Under the Leading Case, if employers are maxi-miners then they will be reluctant to hire greens at any x because they do not have as much experience with greens. This outcome will be the case if we assume, as is plausible, that the "ambiguity set" (cf. (6.4) below) is larger with a worst, worst case scenario, the less experience the employer has with a particular type. If this lack of experience combined with maximin behavior is the cause, then a subsidy for employment of greens will induce experience, which raises the worst case scenario, which will eliminate the resistance of employers if hiring greens at the true $\pi(2,x,\text{green})$ is truly profitable (Brock (2004)).

Notice that improving greens' access to litigation against employers may induce the very maximin behavior that is causing the problem in the first place. This is so because it may increase the size of the "ambiguity set" that employers would face, thus inducing a more negative worst case scenario. More negative worst case scenarios, even if highly improbable, also make minimax regret employers less likely to hire greens. But, at least, minimax regret employers will hire some positive fraction of the stratum (x,green) if the best case $\pi^*(2,x,\text{green})>0$.

Notice also that under Case 1, if employers have less experience at higher levels of x, i.e. the

(6.2) "experience gap", call it G(x),

increases with x (e.g. the worst case $\pi_*(2,x,\text{green})$ becomes more negative as x increases) then the callback rate would fall as x increases if employers are maxi-miners. For minimax regret employers the callback rate would fall only if the ratio of the best case $\pi^*(x,\text{green})$ to the width of the ambiguity interval,

$[\pi_*(x,\text{green}),\pi^*(x,\text{green})]$,

falls with x.

This possibility could be "tested" by modifying the experiment. Suppose CV's are sent out for a new type, violet, where everyone knows that

$[\pi_*(x,\text{violet}),\pi^*(x,\text{violet})]$,

has a worst worst case scenario because employers have even less experience with violets than with greens. I.e. the experience gap is maximal for violets. If the callback rate for violets is higher than greens, the case for discrimination against greens would be strengthened. Furthermore, in the field, if one observed a company with lots of green employees who had been there for quite some time and who were fired after a new president took over, one would also be suspicious.

Manski (2004a) has focused attention on the objective of minimax regret in contrast to maximin because, among other things, as we have seen in this article, maximin is too rigid in its reponse to additional data. We can ask at least two questions, for example, how might we use hypothetical experiments like those above to reveal evidence of minimax regret behavior on the part of employer in contrast to maximin behavior? Will minimax regret "look like" discriminatory behavior as does maximin?

We can directly apply our work above to compute minimax regret. The minimax regret employer will employ a positive fraction, d(x,w) if

(6.3) $0<d(x,w)=\pi^*(x,w)/[(-\pi_*(x,w))+\pi^*(x,w)]$,

even though the worst case, $\pi_*(x,w)<0$, where the maximin employer hires no one from stratum, (x,w). Of course the minimax regret employer is still vulnerable to being unduly influenced by a perceived-to-be very large, worst case belief, $(-\pi_*(x,w))$, for stratum (x,w) even though the probability of this value is tiny under any plausible set of priors (Epstein and Schneider (2002)) that the employer may have.

The idea of using knowledge gleaned from data and other sources to "trim" the set of "admissible" values of π which was used by Brock, Durlauf, and West (2004) in the context of monetary policy under model uncertainties might be used here to improve the performance of maximin and minimax regret or any criterion that is "unduly" impacted by "extreme" and "implausible" values of π.

In any event, the "experiments" that the minimax regret employer conducts under (6.3) should cause it to learn much faster than the maximin employer. E.g. in population, the ambiguity set,

(6.4) $\mathbb{A}:=[\pi_*(x,w),\pi^*(x,w)]$,

shrinks to a point for any stratum where $\pi^*(x,w)>0$, even though $\pi_*(x,w)$ may be very negative, after only one period for the minimax regret employer (Brock (2004)). In contrast, under an adaptive learning scheme studied by Brock (2004), the maximin learner gingerly steps into the employee pool, (x,w), and creeps towards reducing the ambiguity set, $[\pi_*(x,w),\pi^*(x,w)]$, to a point at a rate depending upon the assumed bound on a Lipschitz constant determined by $\pi(2,x)$.

7. SUMMARY, CONCLUSIONS, SUGGESTIONS FOR FUTURE RESEARCH AND APPLICATIONS

This paper has attempted to connect several areas of research that appear unconnected at first thought: (i) Treatment effects and profiling analysis under partially identified structures; (ii) Zero sum game theory, and (iii) Robustification analysis. We compared maximin objectives with minimax regret objectives in two settings. The first setting, "bottoms up," made minimal assumptions on unknown structures. Hence ambiguity sets were wide. The second setting, "top down," considered narrower ambiguity sets centered around a "baseline". We applied this analysis to a lender's problem and showed that an employer's problem had similar analytic structure.

A general conclusion is that the minimax regret objective has nice properties in comparison to the maximin objective for decision making under ambiguity. But the "lotteries" conducted under the minimax regret objective may raise issues of perceived fairness. In general, decision making under ambiguity raises issues of perceived discrimination. Another property of the minimax regret objective that seems under appreciated is that the "lotteries" it conducts can lead to rapid learning of important unknown features of populations and, thus, reduce the size of ambiguity sets quickly relative to some other objectives of decision making under ambiguity such as maximin (Brock (2004)).

Turn now to discussion of future research and potential applications and extensions.

First, we have said nothing about issues raised by the reality of finite samples. Since our prototypical problem, e.g. the lender's problem, is essentially a special case of Manski (2004a), the natural extension would be to adapt his discussion of "empirical success strategies" and the "analogy principle" (which replaces any population object with its sample counterpart) to produce a discussion of sampling issues for any population discussion in our work above. Stoye (2004) has made progress in this direction. Coupling this extension with adaptive learning, especially a comparison of adaptive learning under minimax regret objectives compared to maximin objectives, would be a good topic for future research.

Second, Manski's paper for this Feature, (2004c), studies the problem of a planner who wishes to minimize cost of crime control under partially identified offense functions, the set of which is constrained by data on observed policing search rates and realized offense rates, and a monotonicity assumption on the response of the offense rate to increased police searches. It would be interesting to study adaptive learning by police in his framework, but parallel to our comparison between maximin learning and minimax regret learning.

Third, the lender's problem is close to the classical literature on markets under screening, adverse selection, etc. which was plagued by nonexistence problems. It would be interesting to extend the work above to equilibrium contexts and to discuss empirical issues raised by the interaction of partially identified structures with equilibration. Nonexistence problems could be potentially dealt with by using noisy discrete choice and taking the limit as the "intensity of choice" goes to zero. This idea is closely related to ideas in game theory under the label of "quantal response" games as well as ideas in dynamical systems theory under the label of "natural" limiting measures.

Fourth, the lender's problem could be extended to more realistic settings, e.g., risk averse lenders, nonseparability of utilities across covariates, and relaxation of the size restriction on loans. Some work on this is done in an "unabridged" version of this paper (Brock (2004)).

Fifth, a full treatment requires adjustment for phenomena such as selection bias. Brock (2004) contains some very preliminary results that show connections between different forms of selection bias in treatment choice and ideas from monotone comparative statics such as increasing differences assumptions as well as potential usefulness of assumptions on stochastic dominance. These issues arise when one must use observational data to inform treatment choice and a choice <u>must</u> be made. For example such cases are important when one does not have access to randomized experiments, or the "randomized experiments" are "contaminated" by compliance, attrition, subjects "jockeying" for what they perceive is the "best" treatment, and other problems in reality. See Manski (2004b, 2005, Chapter 2) for results on problems where the planner observes outcomes from an unknown selection process when treatments have been administered in the past. Some work in the context of this paper is done in Brock (2004).

Sixth, it is important to generalize this work to include possible dependence of the treatment effect of treatment t on stratum x on impacts of treatments received by others. In this case we will lose the analytical tractibility of the separability of utilities across x's.

Seventh, in the "top down" approach we only considered one type of "local neighborhood", where the weight on the squared deviation around a "baseline", was unity for each x, for the adversarial agent. We could consider various types of weights to place upon each x as well as neighorhoods generated by different choices of norms, for example, the Euclidean norm and the maximum norm treated above. More on all of this, as well as extended treatments of issues raised by selection bias, unknown covariances across covariates, risk averse objective functions (e.g. risk averse lenders), is in Brock (2004).

Footnotes

1. I thank Chuck Manski, Elie Tamer, and a referee for very useful comments on this paper. I especially thank the referee for expositional assistance and Chuck Manski for his help during the writing of this paper. I thank the NSF and the Vilas Trust for essential financial support. None of the above are responsible for errors, shortcomings, or views expressed in this paper.

REFERENCES

Brock, W., (2004), "Profiling problems with partially identified structure," SSRI W.P. # 2004-21, Department of Economics, University of Wisconsin, Madison.

Brock, W., Durlauf, S., (2004), "Elements of a theory of design limits for Optimal Policy," THE MANCHESTER SCHOOL, 72, Supplement 2, 1-18.

Brock, W., Durlauf, S., West, K., (2003), "Policy evaluation in uncertain economic environments," BROOKINGS PAPERS ON ECONOMIC ACTIVITY, (1): 235-322.

Brock, W., Durlauf, S., West, K., (2004), "Model uncertainty and policy evaluation: Some theory and empirics," Department of Economics, The University of Wisconsin, Madison. JOURNAL OF ECONOMETRICS (in press).

Durlauf, S., (2005), "Assessing Racial Profiling," ECONOMIC JOURNAL (forthcoming) this issue.

Epstein, L., Schneider, M., (2003), "Recursive Multiple Priors," JOURNAL OF ECONOMIC THEORY, 113, 1-31.

Karlin, S., (1959), MATHEMATICAL METHODS AND THEORY IN GAMES, PROGRAMMING AND ECONOMICS, VOLUME I, Addison-Wesley: Reading, Massachusetts.

Manski, C., (2004a), "Statistical Treatment Rules for Heterogeneous Populations," ECONOMETRICA, Vol. 72, No. 4, July, 1221-1246.

Manski, C., (2004b), "Minimax-Regret Treatment Choice with Missing Outcome Data," Department of Economics, Northwestern University, JOURNAL OF ECONOMETRICS (forthcoming).

Manski, C., (2004c),"Search Profiling with Partial Knowledge of Deterrence," Department of Economics, Northwestern University (forthcoming, ECONOMIC JOURNAL, this issue).

Manski, C., (2003), PARTIAL IDENTIFICATION OF PROBABILITY DISTRIBUTIONS, New York: Springer.

Manski, C., (2005), SOCIAL CHOICE WITH PARTIAL KNOWLEDGE OF TREATMENT RESPONSE, Princeton: Princeton University Press.

Stoye, J., (2004), "Minimax Regret Treatment Choice with Finite Samples," Department of Economics, Northwestern University.