Bigots or Bayesians? A model of profiling in law enforcement

Draft 2

September 2005

Linda Welling* and Frances Woolley**,

*Linda Welling: Department of Economics, University of Victoria, PO Box 1700, Victoria, BC Canada V8W 2Y2, email: lwelling@uvic.ca , telephone: (250) 721 8546, FAX: (250) 721 6214

**Frances Woolley: Department of Economics, Carleton University, 1125 Colonel By Drive, Ottawa, Canada K1S 5B6 Phone: 613 520-2600 x 3756, Fax: 613 520-3906, e-mail: frances_woolley@carleton.ca

The first draft of this paper was presented at the meetings of the Canadian Economics Association, Ottawa, Ontario, May 29-June 1, 2003

Abstract:

It is prohibitively costly to investigate every crime. Given limited resources, police and other enforcement agencies may rationally use "statistical discrimination" or “profiling” to minimize search costs. However, such strategies may perpetuate inequities across recognizable groups within the population if different groups have different levels of criminal behaviour and enforcement agencies are imperfectly informed about these levels. We develop a simple game-theoretic model demonstrating this, and use the framework to discuss recent controversies over "racial profiling" and differential sentencing.

Bigots or Bayesians? A model of profiling in law enforcement

Introduction:

Recently there has been a substantial increase in economic analysis of 'racial profiling' (see, for example, Persico (2002), Persico and Todd (2005), Leung, Woolley et al (2002), Farmer, A. and D. Terrell (2001)). Common concerns in this literature include the trade-off between efficiency and equity. Profiling may increase the efficiency of police efforts. For example, suppose that Europeans like strong, ripe, unpasteurized cheese – a restricted important in many countries -- more than members of other ethnic groups do. Then if customs officials' aim is to seize the maximum number of illegally imported dairy products, then intuitively one would expect European-looking people to be stopped and searched more often than members of other ethnic groups. Yet profiling can come with a cost in terms of equity considerations. For example, it might be considered unfair to search one group more intensely than another. Other issues addressed in the literature include the optimal use of resources to maximize convictions, or the best way of minimizing conviction of the innocent.

The existing literature is, however, for the most part, static. For example, Alexeev and Leitzel (2002) consider optimization decisions of individual police officers maximizing the number of convictions. Bunzel and Marcoul (2003) provide a dynamic model in which racially unbiased police who base their priors about criminal propensities on the existing prison population can generate long run prison populations which differ from the racial distribution of the criminal population. In their model, crime rates are exogenous: there is no feedback from police activity to criminal activities.

A related critique is that the literature abstracts from institutional and historical considerations. Discrimination has existed in the past. Obvious examples include slavery of African-Americans in the US and elsewhere, or laws restricting the economic and political participation of Aboriginals in Canada and elsewhere. Past history might plausibly affect present outcomes if beliefs about victimization or criminality are persistent, that is, people are “bigots”. Can beliefs persist even when people are Bayesians, updating beliefs rationally in response to changes in evidence? Can psychological factors, such as confirmatory bias (Rabin and Schrag (1999)), interfere with Bayesian updating? What is the role of factors such as differential search and enforcement costs?

In this paper we address these questions in an extremely simple framework. We consider a world in which there are three types of agents, group R (reds), group G (greens), and the police. Reds have, on average, lower returns to legal activities than greens, but the two group are equally productive (have the same gross return) in criminal work. Each individual’s group membership is common knowledge. Police know the distribution of returns to legal work within each group, as well as the return to crime, but cannot observe the returns to legal activity for individual group members. Individuals must choose whether to engage in legal or criminal activities, a choice that depends upon the gross returns to each type of activity, the punishment or fine associated with crime, and the perceived probability of being caught.

Individuals’ assessments of how likely they are to be caught depend upon what they believe about police behaviour. One novel feature of our model is the structure of these beliefs. We consider various possibilities: people believe police profile; people believe police are neutral; and one group believes police profile while the other does not. Individuals’ beliefs about police behaviour influence their decisions. What generates interesting results, however, is that police do not know what people believe.

In the bad old days of the 1970's, it was possible to examine a situation where agents on one side of a market consistently mis-estimated the value of a choice variable of another agent. **(Ref: warranty papers where consumers had wrong probability of failure.)** Such assumptions do not fit easily into our models of "rational agents", especially when such misperceptions affect the division of any gains from trade. Failing any obvious alternative, the discipline shifted to examining long-run equilibria, where perceptions must be correct, on average. The logical powers required in such equilibria are sometimes astounding.

More recently, work by behavioural economists such as Matt Rabin and Robert Frank have brought to our attention evidence from psychology that suggests rationality may be only partial. Agents in their models are rational, but not perfect: they make mistakes, and some of these mistakes are systematic. To this end, in this paper we explore some consequences of limited learning ability on public perceptions of police activity. This allows us to draw a distinction between Bayesian behaviour – profiling based on some sort of rational, conviction maximizing or crime reducing allocation of police resources – and bigotry, a targeting of police resources towards one group that cannot be explained by optimizing behaviour on the part of police.

III. The model

<u>The basic model: overview</u>

Suppose police are agents of a benevolent state, so their goals are punishing and discouraging criminal acts. They want to distinguish between criminals and non-criminals, but this is costly. They can divide the population at large into two groups: reds (R) and greens (G). There are 2N citizens in total, with fraction $\alpha \in (0,1)]$ belonging to group R. We assume that N is sufficiently large that population characteristics carry over to the population sub-samples subject to police investigation. The groups are identified by a criterion can be observed minimal cost (set to zero for convenience) for example, colour, approximate age, or sex. If police engage in profiling, in the first instance their treatment of any individual is determined by that individual's group membership.

Police have limited resources (budget B>0), which are not sufficient to investigate fully every crime or every individual. Any attempt at identifying criminals will therefore involve some tradeoffs. The technology for identifying criminal activity is ‘scanning’. Scanning correctly identifies a criminal as such with probability $p_s \in (0.5,1)$, and correctly identifies a non-criminal as innocent with probability $m_s \in (0.5,1)$. Associated with this technology is a constant unit (resource) cost, $d > 0$.[1]

Police have fixed resources, B, which they allocate between the two groups subject to the budget constraint $B_g + B_r = B$. This resource allocation decision determines the number of individuals in each group who will be scanned, $n^i = \frac{B_i}{d}, i = r, g$. The fraction of individuals scanned in each group is $\theta_i \equiv \frac{n_s^i}{N_i}, i = r, g$. The assumption that $B < \min d\{N_r, N_g\}$ ensures it is not feasible to scan all members of either group.

Optimization problems for reds and greens

Individuals have a choice between allocating all of their time to legal activities or all of their time to crime. The return to legal activities is higher for group G (and as a result, we will find that reds get stopped and scanned more often, hence red for stop). To be more precise, we assume that the distribution of returns to legal activities for the two groups are uniform on $[\underline{w}^i, 1]$, with $\underline{w}^r = 0, \underline{w}^g = \beta + \underline{w}^r = \beta$, and $\beta \in [0,1)$. The two distributions have the same upper bound, but the lower bound of the Green distribution is higher. The gross return from crime, y, is assumed to be the same for all individuals, and β <y<1. Diagram 1 illustrates the supports of the distributions.

Diagram 1

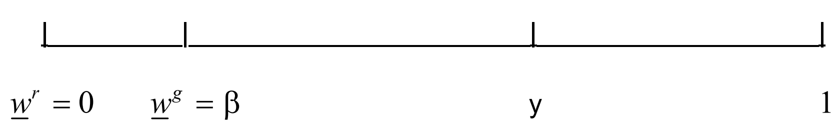

[1] We consider heterogeneous screening costs in an extension.

Individuals are assumed to be risk-neutral; that is, $U_j = w_j$, where w_j is the wealth of individual j. Individuals therefore choose to maximize expected wealth. If a person chooses to engage in legal activities, expected wealth is simply w_j less any expected cost associated with inspection (c^i) and "wrongful conviction" (f) for criminal activities. Expected wealth when a person engages in criminal activities is the gross return to crime less the expected value of the costs of being investigated and caught.

Individuals bear a cost if they are scanned: this cost is denoted by c^i for individual of group $i = r, g$. Given possible screening, an individual of group $i = r, g$ who commits no crime receives one of three payoffs (see Diagram 2) :

a) if not scanned, receives w_j for certain - this occurs with probability $(1-\theta_i)$;

b) if scanned and correctly identified, receives $w_j - c^i$ - with probability $\theta_i m$;

c) if scanned and incorrectly identified, receives $w_j - c^i - f$ - with probability $\theta_i (1-m)$.

Thus, an individual of group $i = r, g$ who chooses *not to commit a crime* has expected payoff

$$E\Pi_n^{\ i} \equiv w_j - \theta_i [c^i + (1-m) f] \tag{0.1}$$

Notice that the expression in [ ] is the sum of expected (psychic) costs and the fine this individual incurs if mistakenly identified as a criminal.

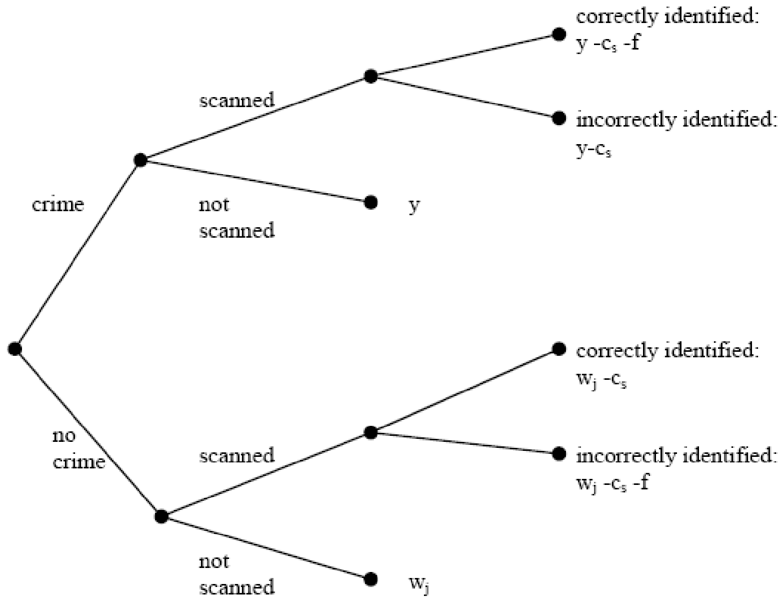

The outcomes, and their associated probabilities, for an individual who chooses to commit a crime can be derived analogously. The only significant differences in the derivations are that the gross return from crime is *y* for all individuals, and payment of the fine *f* results from correct (rather than erroneous) identification with probability p_s. The expected payoff for an individual criminal is therefore

$$E\Pi_y{}^i \equiv y - \theta_i[c^i + pf] \qquad (0.2)$$

Thus, crime will pay for an individual belonging to group $i = r, g$ if and only if

$$w_j \leq \bar{w}^i \equiv y - \theta_i[f(p+m-1)] < y \tag{0.3}$$

The term in [ ] is the expected fine facing an individual who commits a crime, based on the absolute value of the fine, the probability of being scanned, and the probabilities of Type 1 and Type 2 errors in the screening technology.[2] Each individual, in determining the value of crime versus legal activity, weights this expected fine by θ_i, that individual's probability of being screened, which is based on group affiliation.

The last inequality holds providing $p+m>1$. If $p+m=1$, tossing a fair coin to determine whom to prosecute would be equally productive; this inequality stipulates that the scanning technology must be more accurate than a coin toss. The cutoff level of wealth, $\bar{w}^i$, is increasing in y, and strictly decreasing in both the ultimate punishment, f, and the probability of detection. The cost to an individual of being scanned does not appear, since the probability of being scanned is independent of any action on the part of an individual, and hence this cost has no deterrent value.

All individuals in group $i = r, g$ with initial wealth less than $\bar{w}^i$ choose to engage in crime, and none of those with higher wealth will do so. Given our assumed wealth distributions, the proportion of individuals of group R who commit crimes, defined above as π_r, is equal to $\bar{w}_r$, and the absolute number of individuals in group R who commit crimes is equal to $2\alpha\bar{w}^r N$. The corresponding values for group G are $\pi_g = \dfrac{(\bar{w}^g - \beta)}{1-\beta}$, and $\dfrac{2(1-\alpha)N(\bar{w}^g - \beta)}{1-\beta}$, respectively, hence the expected number of criminals, in any period, is

$$\pi_r N_r + \pi_g N_g = 2(\alpha\pi_r + (1-\alpha)\pi_g)N.$$

[2] Type 1 errors occur when we reject the null hypothesis when it is in fact true (false positive or wrongful conviction); Type 2 error occurs when we accept the null hypothesis when it is in fact false (false negative or allowing the guilty to go free). The probability of a Type 1 error is 1-m_s, the probability of Type 2 is 1-p_s

The overall population size N affects the absolute number of crimes committed by each group, both directly and, given a fixed police budget, through its effect on the proportion of people scanned. Differences in relative group size (α) affect the proportion of total crimes committed by each group. When a group has relatively better legal work opportunities members of that group commit relatively less crime: more precisely, when β, group G's lower bound on returns to legal work, increases, the number of group G members for whom criminal activities are more attractive than legal activities will decrease. Therefore variations in β would be expected, all else being equal, to affect the fraction of individuals from group G who commit crimes, as well as the ratio of (criminals from group R) to (criminals from group G).

Law enforcement technology and the resource allocation decisions of police affect both absolute and relative crime levels through their effects on $\overline{w}$: when scanning is more effective, fines increase or more police resources are allocated to a particular group, $\overline{w}$ falls, and with it the proportion of criminals in each group. Law enforcement technology will also affect the number of crimes detected. Differences in the absolute or relative number of total crimes committed by individuals from the two groups can arise due to differences in group size or the distributions of outside opportunities. Differences in the amount of criminal activity detected may reflect these same factors, or law enforcement technology or the resource allocation decisions of police.

This brings us to a crucial point: an individual's decision to engage in criminal behavior is contingent upon what his or her beliefs about police behavior and law enforcement technology -- what are the odds that I will get caught? To discover the possible equilibrium allocation of resources, we need to describe the police's optimal resource allocation decision, and fully explore the interaction between police and (potential) criminal's optimization decisions.

<u>Full information</u>

In this section we present solutions for our model when there is no uncertainty on the part of any of the three types of agents. This is the type of equilibrium that we have come to expect in economic models, following the rational expectations and game theory revolutions. Below we consider a Nash equilibrium, where each agent takes as given the choices of others; the case where police take into account the deterrent effects of their choices is in Appendix A. This case provides intuition as to the workings of this basic model, and serves as a benchmark against which we can compare later results and the results of other models.

Case 1: Nash equilibrium

In the Nash equilibrium, police take the decisions of the two groups as given, while individuals take as given the police resource allocation. We have solved for the individual's optimization decision above. Here we re-state that solution, determine the optimal police response and the Nash equilibria. Criminal subsets of groups R and G are as determined as follows: each individual takes as given their own returns to legal and illegal activities and the police resources allocated to their group, and optimally decides whether or not to engage in crime. There will be a cutoff level of w such that all individuals with returns to legal activity above that level will choose legal work, and all below will choose crime. Then, the criminal population will consist of the following:

Group G: $2(1-\alpha)\pi_g N$, individuals from returns to legal activity in the range $[\beta, \bar{w}_2]$, where

$$\bar{w}_2 \equiv y - \frac{B - B_r}{2(1-\alpha)Nd}(p+m-1)f$$

(*Notice that there will be no criminals from group G if $\beta > \bar{w}_2$.)

Group R: $2\alpha\pi_r N$, individuals with returns to legal work in the range $[0, \bar{w}_1]$, where

$$\bar{w}_1 \equiv y - \frac{B_r}{2\alpha Nd}(p+m-1)f$$

Where $\pi_i, i = r, g$ denotes the fraction of the relevant group engaged in crime.

Taking the absolute number of criminals from each group as given[3], police allocate resources to maximize their objective function. Initially we assume that police aim to solve crimes, however there are a number of other plausible alternative objective functions, including minimizing the number of crimes committed or maximizing fine revenue:

Criterion I: maximize the expected number of correct convictions.

In this case the police objective function is

$p(\pi_r \frac{B_r}{d} + \pi_g \frac{B - B_r}{d}) = [B_r(\pi_r - \pi_g) + \pi_g B]p / d$. Given the linearity in the problem, the solution (the optimal choice of $B_r \in [0, B]$) is trivial. If criminal activity is chosen by different proportions of the two populations, all resources should be devoted to the group with the larger fraction of criminals. If $\pi_r = \pi_g$ then, *taking individual behaviour as given*, resources can be divided in any way between the two groups. If the same proportion of each population chooses crime, then the two populations are essentially identical, and screening individuals on the basis of group affiliation is no more efficient than ignoring affiliation.

Imposing consistency across the decisions of all agents results in two types of equilibria[4]

a) Targeting group R: $B_r = B, B_g = 0$. This will be the equilibrium if a scan of a group R individual has a higher probability of correctly convicting a criminal than that of an individual from group G. This occurs when the average incomes of the two groups are sufficiently different that only a small fraction of group G find crime attractive, even when it is certain that they will not be stopped (so the expected penalty for crime for individuals from this group is zero). Substituting in for the criminal populations yields this as an optimal policy when

$$\pi_g^{\max} = \frac{(y - \beta)}{1 - \beta} \leq [y - \frac{B}{2\alpha Nd}(p + m - 1)f] = y - \overline{F}_r = \pi_r^{\min}$$

[3] The alternative is that police act as a Stackelberg leader, taking as given the responses of individuals in the two groups to police resource allocation.
[4] Not multiple equilibria, since which type occurs depends on parameter values.

Targeting group R is most likely when when police have few resources (low ratio of B/N), scanning is costly (high d) or inaccurate (high p, m) and the potential fine, f, is small. Taken together, these group of parameters capture the effective resources available to police. The fewer resources available, the greater the likelihood that all available resources will be concentrated on the group fewer attractive alternatives to criminal behaviour. When group R is relatively large (higher α) it is more likely to be the exclusive target of police investigations. Also, as the gap between the rewards to legal work available to the two groups, β, widens, group R again will tend to be targeted more often.

The quotient on the LHS of this inequality gives the fraction of group G who choose to engage in crime when there is no penalty, so this is the *maximum* possible number of criminals from group G. The expression on the RHS gives the number of individuals from group R who choose to commit crime when they correctly anticipate that *all* available police resources will be devoted to them; this then is the *minimum* fraction choosing crime from group R. This inequality specifies that targeting R is optimal when there are proportionately more criminals from group R than from group G. This inequality can be expressed in terms of a lower bound on returns to legal activity for members of group G, in terms of the expected fine facing the R group:

$$\beta \geq \frac{\bar{F}_r}{1 - y + \bar{F}_r} > 0$$

where $\bar{F}_r \equiv \frac{B(p_s + m_s - 1)f}{2\alpha Nd}$ is the maximum expected fine an individual criminal from group R can face, that is, the fine he or she would expect to face if in the group to which all police resources are allocated.[5]

If $\beta < \frac{\bar{F}_r}{1 - y + \bar{F}_r}$, it cannot be optimal for the police to target group R, since switching some resources to group G would increase the expected number of convictions.

b) Targeting group G: $B_r = 0, B_g = B$: this is never a unique equilibrium, because it requires $\beta \leq -\frac{\overline{F}_g}{1-y} < 0$. [Here, $\overline{F}_g = \frac{\alpha \overline{F}_r}{(1-\alpha)}$ is the maximum expected fine for an individual criminal from group G. Notice that if $\alpha = 1-\alpha = 0.5$, then $\overline{F}_g = \overline{F}_r$, whereas $\overline{F}_g < \overline{F}_r$ if $\alpha < 0.5$.] Since we have assumed a lower bound of zero on the returns to legal activity for group G individuals, this will never occur.

Thus, targeting group R will be the optimal solution if the groups are sufficiently different. If the two groups are relatively similar, a different solution arises:

c) Scanning both groups: If the conditions above are not satisfied, so $\beta \in (0, \frac{\overline{F}_r}{1-y+\overline{F}_r})$ targeting one group to the exclusion of the other is not an equilibrium. The only possible equilibrium occurs when $\pi_g = \pi_r$. In this case, screening one fewer person from group R and one more from group G will not change the likelihood of convicting a criminal, so the police are indifferent *at the margin* to the allocation of resources. One might think there would be a *prima facie* case in this situation for dividing police resources equally between the two groups. However, individuals base their choices on the (assumed) allocation of police effort. Corresponding to each value of β is a unique resource allocation, B_r^*,[6] for which the following condition holds:

$$\frac{y - \frac{(B-B_r)}{B}\overline{F}_g - \beta}{1-\beta} = y - \frac{B_r}{B}\overline{F}_r,$$

---

[5] This cost is increasing in total resources per capita, the fine, and the efficiency of the scanning technology (higher p_s, lower d). Notice that since there is a possibility that an innocent person might be found guilty, this fine is really the incremental expected fine from committing a crime.
[6] This assumes a fixed B

Rearranging this equation, and solving for the optimal resources devoted to group R, shows that this to be fraction γ

$$B_r = \gamma B, \quad \text{where } \gamma \equiv \frac{\beta(1-y)+\bar{F}_g}{\bar{F}_g + \bar{F}_r(1-\beta)} = \frac{\beta(1-y)+\bar{F}_g}{\bar{F}_g + \bar{F}_r(1-\beta)}$$

$$= [\alpha + \frac{\beta(1-\alpha)(1-y)}{\bar{F}r}]/[1-\beta(1-\alpha)]$$

$$\geq \alpha \;\; \forall \beta \geq 0$$

To interpret this, suppose first that the distributions of returns to legal activities are identical, so β = 0. Simple substitution shows that when β=0, $\gamma = \alpha$: when the populations have identical income distributions, the apportioning of the police budget should reflect the relative proportions of the two groups in the population :

$\frac{B_r}{B - B_r} = \frac{\alpha}{1-\alpha}$. If α=0.5, there will be equal numbers of criminals from both groups when half the available resources are devoted to each group. Whereas if α < (>) 0.5, equal numbers of criminals from each group can be the outcome only if the police devote less (more) than half of the resources to the smaller group.

It is trivial to show that γ , the share of the police budget allocated to group R, is increasing in {N,d, β } and decreasing in {B,f, p, m }.

The more interesting case, however, is when group R (G) has a relatively higher return to illegal (legal) activity than the other group, that is, $0 < \beta < y$. Since γ is increasing in β , when α =0.5 more than half the available resources must be devoted to group R. In all cases, the proportion of resources devoted to the less fortunate group exceeds their proportion in the overall population.

γ attains its maximum value of unity when $\beta = \dfrac{\bar{F}_r}{1 - y + \bar{F}_r}$. Recall that

$\bar{F}_r \equiv \dfrac{B}{2\alpha Nd}(p + m - 1)f$; it then follows that the greater the fine, the more equal will be the resources devoted to the two groups in the mixed equilibrium.

One interpretation of this equilibrium is that by diverting more resources away from population G to population R, police are *encouraging* criminal behaviour on the part of those members of population G with lower potential income. The more resources are directed at R, with a given total budget, the greater the difference in the expected return from crime between the marginal criminal from G and the marginal criminal from R. This means the criminal population is drawn from higher (potential) income groups in population G than R.

<u>Discussion</u>

The equilibria above are illustrated in diagram _. The knife-edge equilibrium is interesting because common behaviour is induced/supported across the different populations by what might be interpreted as discriminatory behaviour on the part of the police. When the groups are relatively similar but not identical, allocating resources equally between the two groups will result in fewer criminals from group G than from group R. Switching more resources to group R will balance the number of criminals across the two groups. If welfare is improved by equal outcomes, then this form of discrimination will be preferred to "non-discriminatory" behaviour on the part of police.

Alternative Specifications of the Police Objective Function

The results presented here are robust to at least some alternative specifications of police objectives. It could be argued, for example, that the appropriate objective function is the total number of convictions, as opposed to the number of correct convictions. This would be a plausible objective function is police were unable to differentiate between “true positives” and “false positives”. An equivalent way of thinking about this objective is that the police aim to maximize fine revenue – a suggestion which finds some empirical support [REFERENCE]. In this case, the police objective function (criterion 1 above) would be replaced with criterion 2:

Criterion 2: maximize the total number of convictions (correct or incorrect).

Formally, maximizing the number of criminals correctly convicted plus the number of people engaged in legal activities but incorrectly identified as criminals gives the objective function:

$$p(\pi_r \frac{B_r}{d} + \pi_g \frac{B - B_r}{d}) + (1-m)((1-\pi_r)\frac{B_r}{d} + (1-\pi_g)\frac{B - B_r}{d})$$

$$= [(1-m) + (p+m-1)\pi_r]\frac{B_r}{d} + [(1-m) + (p+m-1)\pi_g]\frac{B - B_r}{d}$$

$$= \{(p+m-1)(\pi_r - \pi_g)B_r + [(1-m) + (p+m-1)\pi_g]B\}/d$$

The objective function above has, however, the same linearity properties as the initial objective function and, therefore, has the same two equilibria. If $\pi_r^{min} > \pi_g^{max}$ the objective function is maximized by devoting all effort to group r; when this inequality does not hold, any allocation of effort will maximize the police objective function, but consistency of police effort allocations with individual decisions requires the particular allocation yielding $\pi_r = \pi_g$, as in the solution described above.

A third possible specification of the police objective function is that police attempt to minimize the portion of the population engaging in criminal activity, that is, minimize $\pi_r + \pi_g$. If, as suggested above, the police take these ratios as given, there is no effort allocation which yields the best solution here. When police take into account individuals' responses to police effort, this becomes a feasible objective. We consider this case in the appendix, where we consider Stackelberg behaviour on the part of police.

A fourth possible specification of the police objective function allows for differential screening costs across the two groups. Screening costs may differ because of the psychic costs or benefits of screening different groups – for example, Knowles, Persico and Todd (2001) suggest that police screen women for possession of drugs more often than is optimal in terms of maximizing the number of drug seizures because of the lower costs/higher benefits to police of searching women. More generally, a person who has ready access to legal counsel may be more costly to interrogate and more difficult to extract a 'confession' from than someone without education or access to legal advice. The possibility of different screening costs alters the police objective function we have:

Criterion 4: maximize convictions allowing for differential search costs

With the associated objective function:

$$p_s(\pi_r \frac{B_r}{d_r} + \pi_g \frac{B - B_r}{d_g})$$

where d_i is the screening costs associated with members of group i. Introducing differential screening costs means that the formulae for the maximum fines faced by individuals in the two groups are further differentiated: with some abuse of notation, for what follows we redefine $\overline{F}_r = \frac{B(p+m-1)f}{2\alpha d_r N}$ and $\overline{F}_g = \frac{B(p+m-1)f}{2(1-\alpha)d_g N} = \frac{d_r}{d_g}\frac{\alpha}{(1-\alpha)}\overline{F}_r$. Let the ratio of screening costs be denoted by δ, so $\delta = \frac{d_r}{d_g}$. Depending on how δ varies relative to the proportions of the two groups in the population, the maximum fines for the two groups may be more or less similar than those with a common screening cost.

Again the linearity of the objective function ensures that the problem has the same two types of solution: targeting and partial pooling. The targeting equilibrium will prevail when $\frac{\pi_r^{\min}}{d_r} > \frac{\pi_g^{\max}}{d_g}$ or $\frac{\pi_r^{\min}}{\pi_g^{\max}} > \delta$: that is, when the parameters are such the even when all resources are devoted to group R, sufficiently few individuals from group G choose to engage in crime that, given the high cost of screening group G individuals and the given police budget, it is more efficient to ignore criminals from group G. When this inequality does not hold, expending some resources on group G will be optimal. The allocation of police resources affects the proportion of individuals from each group who chose to commit crimes (π_i), and this optimal allocation is such that

$\frac{\pi_r}{d_r} = \frac{\pi_g}{d_g}$ or $\frac{\pi_r}{\pi_g} = \delta$: that is, when the relative crime rates are equal to the relative scanning costs. Notice that the targeting equilibrium here is, in some sense, less appealing. Previously the police were targeting all resources to group R because group R had (relatively) more criminals; in this equilibrium police may target all resources to group R even if group G has more criminals, if the cost of scanning members of group G are sufficiently high.

As before, characterizing the mixed equilibrium is slightly more complicated. It will be an equilibrium strategy for police to screen both groups when the budget allocation (B_r) satisfies

$$\frac{y - \frac{(B - B_r)}{B}\overline{F}_g - \beta}{(1-\beta)} = y - \frac{B_r}{B}\overline{F}_r$$

Solving this yields $B_r = \gamma_2 B, \quad \text{where } \gamma_2 \equiv \frac{\beta(1-\alpha)(1-y) + \alpha\delta\overline{F}_r}{[(1-\beta)(1-\alpha) + \alpha\delta]\overline{F}_r}$; this expression collapses to the previous result when $\delta = 1$ (so $d_r = d_g = d$). Notice that, as before, an increase in β leads to more resources being devoted to policing group R. However, the higher the relative cost of scanning group g, the greater the redistribution of resources involved with a given increase in β .

<u>Information problems: learning the system</u>

In this section we relax the assumption that citizens and police are fully informed about each other's activities, and consider the possibilities of agents learning about the system. In the first type of equilibrium above, where police devote all resources to group R, it is not difficult to believe that citizens might learn actual police behaviour: if only group R individuals are scanned, then only members of this group will be convicted. Any statistics police release on actual numbers of conviction by group affiliation would reveal their policy. Even without official release of such statistics (which are not readily available in many countries) individuals would learn, from media reports or the experience of friends or family, that no G individual is ever convicted.

However, suppose the two groups are not that different, and in a full information world the second type of equilibrium would obtain. In this equilibrium police screen members of both groups, and both groups would appear in statistics on actual convictions. Since police cannot identify criminals without investigation, and the police scanning technology is subject to random error, the actual number of convictions from each group would not, in general, equal the expected number (given by $\pi_i, i = r, g$). Therefore the released statistics would represent a noisy signal of police activity. Would individuals observing these signals be able to distinguish between a police policy that ignored group affiliation and one that devoted more resources to group R? And would police be able to determine the appropriate screening proportions, given the conviction rate of those screened?

Our approach to this is still extremely tentative, and much work remains to be done. The following sets up and discusses this approach.

We assume citizens have priors on police allocation of resources - estimates of the fraction of individuals of each group scanned by police, $\theta_i, i = r, g$.[7] Our story is much more interesting if reds and greens, members of group R and G, have different beliefs about police

---

[7]Recall that $\theta_i \equiv \frac{n_s^i}{N_i} = \frac{E_i}{dN}$, so this parameter measures the fraction of individuals of type i who are scanned by police; we retain the assumption that populations are large enough so the

behavior. There are a variety of scenarios that can generate different beliefs, for example, groups R and G are exposed to different media or live in different neighborhoods. However the scenario that we find convincing is as follows: at time t-1, discrimination exists. Police target group R, regardless of the actual proportion of criminals in group R compared with group G. At time t, however, due to (say) a constitutional change prohibiting discrimination, police announce that they will no longer target group R. We then have a variety of different scenarios: the police announcement is believed by both group R and group G; the police announcement is believed by group G but not group R (plausible in the light of group R's previous negative experiences being convicted by the police and group G's lack of interaction with the police, even more plausible if members of the police are predominantly drawn from group G); finally, the police announcement is believed by neither group, and both believe that the police will continue to discriminate.

More formally, consistent with our assumption that $\beta \geq 0$, we assume individuals believe that police either follow the equilibrium strategy, or "target" group R - that is, devote more resources to group R than the proportion dictated in equilibrium; let the fraction of resources devoted to group R in the latter strategy be $\varphi > \gamma_2$. This generates a range for the belief parameter for each group: for group G, $\theta_g \in [\frac{(1-\varphi)B}{2(1-\alpha)d_g N}, \frac{(1-\gamma_2)B}{2(1-\alpha)d_g N}]$, while for group R, $\theta_r \in [\frac{\gamma_2 B}{2\alpha d_r N}, \frac{\varphi B}{2\alpha d_r N}]$. Let $\rho_t^i \in [0,1], i = r, g$ denote the probability with which an individual from group i believes, at time t, that police are devoting more resources to group R (so are bigots, in the eyes of citizens): that is, Pr $\{\theta_r = \frac{\varphi B}{2\alpha d_r N}\} = \rho_t^i, i = r, g$. Then, with probability (1- ρ_t^i) police act to minimize crime, so Pr $\{\theta_r = \frac{\gamma_2 B}{2\alpha d_r N}\} = (1-\rho_t^i), i = r, g$. Given this notation, and risk neutrality, the expected cost of criminal activity for individuals in group R is $[\rho^r \varphi + (1-\rho^r)\gamma_2]\overline{F}_r$. Similarly, the expected cost of criminal activity for individuals in group G

probability of any given individual being scanned is equal to this fraction for the relevant is $[[\rho^{g}(1-\varphi)+(1-\rho^{g})(1-\gamma_{2})]\overline{F}_{g}$. As before, individuals will rationally choose to engage in crime, or not, depending on which activity maximizes expected utility, given beliefs. We can identify cutoff levels of wealth $w_{r}(\rho^{r})$ and $w_{g}(\rho^{g})$, dependent on beliefs, such that all individuals in that group with lower wealth will choose to become criminals, and all individuals with higher wealth will abstain from crime.

Within the framework above, we have four possible sets of *extreme beliefs*:

i) both groups believe police devote more than the crime minimizing resources to group R , so all believe $B_{r}=\varphi B>\gamma_{2}B$ - we call this bigotry. In this case all individuals believe that: police are neutral, so

$$\{\theta_{g}{}^{e},\theta_{r}{}^{e}\}=\{\frac{(1-\gamma_{2})B}{2(1-\alpha)d_{g}N},\frac{\gamma_{2}B}{2\alpha d_{r}N}\}$$

ii) group R believes police are bigots while group G believes in neutrality, so group R has $\{\theta_{g}{}^{e},\theta_{r}{}^{e}\}=\{\frac{(1-\varphi)B}{2(1-\alpha)d_{g}N},\frac{\varphi B}{2\alpha d_{r}N}\}$ and group G has

$$\{\theta_{g}{}^{e},\theta_{r}{}^{e}\}=\{\frac{(1-\gamma_{2})B}{2(1-\alpha)d_{g}N},\frac{\gamma_{2}B}{2\alpha d_{r}N}\};$$

iii) group R believes in neutrality, group G believes in bigotry: this is the converse of the beliefs in (iii)

Clearly, the latter two sets of beliefs cannot be consistent with actual police activity, since the groups' perceptions of that activity differ.

How does the population of criminals change with these different sets of beliefs? First, notice that for any given $\beta\in[0,\frac{\overline{F}_{r}}{1-y+\overline{F}_{r}}]$ we can divide the population of each group into three subsets:

---

population.

those who always/for all beliefs commit crimes, those who never commit crimes (for any feasible beliefs) and those whose criminal activity depends on beliefs. Thus:

Group		Wealth range	
	Always criminals	Depends on beliefs	Never criminals
R	$w \in [0, w_1]$	$w \in [w_1, w_2]$	$w \in [w_2, 1]$
G	$w \in [\beta, w_2]$	$w \in [w_2, w_3]$	$w \in [w_3, 1]$

with $0 < w_1 < w_2 < w_3 < 1$, and [note: should any of these be weak inequalities?]

w_1 is the cutoff wealth for group R when group R individuals believe $B_r = \varphi B$;

w_2 is the upper cutoff wealth for group R , and the lower cutoff level of wealth for group G, when individuals from both groups believe $B_r = \gamma_2 B$; and

w_3 is the upper cutoff wealth for group G when group G individuals believe $B_r = \varphi B$ (and so $B_g = (1-\varphi)B)$.

Regardless of group affiliation and beliefs, given the parameters of the justice system, all those with wealth in range $w \in [0, w_1]$ will always be criminals, and none of those with wealth $w > w_3$ will ever be.

In all the cases we listed above, the criminal population is drawn from both the red and the green groups. Moreover, optimal police policy calls for screening some individuals from both groups. It follows that, in general, the set of convicted individuals would be comprised of members of both groups. This set is more likely to have a higher representation from group R if both groups put higher weight on police being neutral than if both think it more likely that police are targeting group R.

In their "confirmatory bias" model, Rabin and Schrag (1999) consider an agent subject to error in interpreting signals about the true state of the world. An agent with initial beliefs about the true state of the world receives a sequence of signals correlated with the true state, and this agent updates their priors, according to Bayes rule, upon perception of each signal. Error creeps in because the agent's perception of some signals can be faulty. Rabin and Schrag assume that if the sequence of signals has led the agent to put more weight on state A than on state B, then that agent always correctly perceives signals that bolster current beliefs. However, a signal that would cause a Bayesian observer to attach a lower probability to state B is, with a fixed probability, misperceived as evidence which would increase the prior on state A. Rabin and Schrag show that such bias can result in beliefs that do not converge to a correct interpretation of the world, even in the limit.

The situation we are trying to model in this paper is more complex than the Rabin and Schrag model, because it is not clear that in our model the state of the world is exogenously fixed. In this paper the "state of the world" can be interpreted *either* as the extent to which police activity targets group R, *or* as the extent to which individuals from group R are more likely than those from group G to engage in criminal activity. The first interpretation speaks to public perceptions that police use the correct amount of profiling, while the second concerns the attempt by police to allocate resources fairly and efficiently.

Suppose we follow up on the former interpretation, since it seems more in line with the work above. Then we can consider individuals in group R and group G as the agents. The sequence of signals about police activity is a series of statistics on convictions. Most individuals do not receive or study these statistics *per se*, but rather learn about them from various media. It does not seem much of a stretch to assume that different media present these statistics in different ways. Hence, so long as individuals from different groups access different media,[8] systematic misperception is plausible.

This suggests, then, that it is possible for groups R and G, if they start from different priors over the extent of police profiling, to have beliefs which converge to different opinions. For instance, suppose that initially $\rho^r = 0.5 > \rho^g$. A high ratio of conviction of group R individuals would increase both priors, and might converge to a situation where $\rho^r > 0.5 > \rho^g$. If both groups were subject to confirmatory bias, this could lead to a situation where individuals of group R were convinced that police were targeting them, while individuals in group G were equally convinced that police were ignoring group affiliation. Notice that this combination of priors leads to the smallest possible criminal population, for a given budget, so that it is not clear that the police, acting for society at large, would have any incentive to try and align the priors.

[8] It might be argued that the choice of media is correlated as much, if not more, with income/wealth as with whatever group affiliations R and G represent. However, recall that in this case all we are concerned with is the group of individuals who fall within the $[w_1, w_3]$ range, not those at either extreme.

APPENDIX A:

Case II: Stackelberg equilibrium:

Suppose now that police take into account the deterrent effects of their allocation decision. We retain the assumption that criminals take the police allocation of resources as given, but assume that police can solve individuals' decision problems. Given the goal of maximizing the expected number of convictions, police now choose B_r to maximize

$$p(\pi_r \frac{B_r}{d} + \pi_g \frac{B-B_r}{d}) \text{ subject to } \pi_r = y - \frac{B_r}{B}\bar{F}_r \text{ and } \pi_g = \frac{y - \frac{B-B_r}{B}\bar{F}_g - \beta}{1-\beta}$$

This problem has an interior solution, for some values of the parameters:

$$B_r^* = \gamma_2 B, \text{ where } \gamma_2 \equiv [\alpha + \frac{\beta(1-y)(1-\alpha)}{\bar{F}_r}]/[2(1-\beta(1-\alpha))] = 0.5\gamma =$$

Following the analysis above, notice first that when $\beta = 0$, so the distributions of returns for the two groups are identical, police should allocate half of their resources to each group. As β increases, and the groups become more dissimilar, then γ_2, the fraction of resources allocated to group R, increases, reaching the upper bound of one at $\beta = \frac{\bar{F}_r(2-\alpha)}{(1-\alpha)[1-y+2\bar{F}_r]}$.

Comparison / discussion

As one might expect (and hope!), when police take into account the deterrence effects of their decisions, targeting group R is less likely to be optimal. The upper bound on β in the Stackelberg equilibrium is larger than that in the Nash equilibrium, so that when police take into account the deterrent effects of their activities the range over which the interior equilibrium prevails is larger than the corresponding range for in the Nash equilibrium. If individuals know that more resources are being devoted to group R, then crime is more attractive at the margin to individuals from group G and less attractive to those from group R, all else constant. If the police consider these incentives when they allocate resources, a smaller proportion of resources would be shifted to targeting group R.

Suppose now that the objective is to minimize the proportion of the population engaging in criminal activity, $\alpha\pi_r + (1-\alpha)\pi_g$. We can then substitute for π_i from the individual maximization problems to get the new police objective function

$$\text{minimize } \alpha(y - \frac{B_r}{B}\bar{F}_r) + (1-\alpha)(\frac{y - \frac{B-B_r}{B}\bar{F}_g - \beta}{1-\beta})$$

Rearranging this objective function shows that the coefficient on the variable of interest, B_r / B, is $\frac{\beta\alpha\bar{F}_r}{(1-\beta)} > 0$, which tells us that so long as β>0 minimizing the proportion of the population involved in crime requires devoting all resources to group R , *regardless of the relative sizes of the population.*

Once again, the linearity of the problem dictates that there will, in general, be a corner solution. As long as $\beta > 0$, the marginal deterrent value of an extra unit of effort will be greatest when it is allocated to policing the red group. Only when $\beta = 0$ will police devote any effort to policing the green group.

References:

Alexeev, Michael and Jim Leitzel (2004) "Racial Profiling", available at http://mypage.iu.edu/~malexeev/

Bunzel, Helle and Philippe Marcoul (2003), "Can Racially Unbiased Police Perpetuate Long-Run Discrimination?", Tilburg University, CentER Discussion Paper 2003-16 ISSN 0924-7815

Farmer, A. and D. Terrell (2001), "Crime versus Justice: is there a trade-off?" *Journal of Law and Economics*, vol XLIV, pp 345-365

Knowles, John, Nicola Persico and Petra Todd, 2001. Racial Bias in Motor Vehicle Searches: Theory and Evidence. *Journal of Political Economy* 109(1): 203-229.

Leung, A., F. Woolley, R.E. Tremblan, and F. Vitarol (2002), "Who gets convicted? Statistical discrimination in law enforcement", Carleton Economic Paper CEP 02-03

Persico, Nicolas (2002) “Racial Profiling, Fairness, and the Effectiveness of Policing” *American Economic Review,* vol 92(5), pp. 1472-97.

_____________ and Petra Todd (2005), "Passenger profiling, imperfect screening and airport security" , PIER wp 05-005

Rabin, M. and Schrag, _ (1999), "First impressions matter: a model of confirmatory bias", *Quarterly Journal of Economics*, vol CXIV (1),pp 37-82)

***