Penn Institute for Economic Research
Department of Economics
University of Pennsylvania
3718 Locust Walk
Philadelphia, PA 19104-6297
pier@econ.upenn.edu
http://www.econ.upenn.edu/pier

PIER Working Paper 05-019

"Crime Minimization and Racial Bias: What Can We Learn From Police Search Data?"

by

Jeff Dominitz and John Knowles

http://ssrn.com/abstract=719981

Crime Minimization and Racial Bias: What Can We Learn From Police Search Data?*

Jeff Dominitz†

Heinz School, Carnegie Mellon University

John Knowles‡

Department of Economics, University of Pennsylvania

February 18, 2005

Abstract

Are variations in the success rate of searches by race informative about racial bias if police are motivated by crime minimization rather than success-rate maximization? We show that the basic idea of extracting information from hit rates may still be valid, provided one

*We thank Charles Manski, John Pepper, and Nicola Persico for helpful discussions. We are also grateful for the comments of participants in seminars at Georgetown University, the University of Pennsylvania, and the Conference on Social Interactions in Profiling at Northwestern University.

†dominitz@andrew.cmu.edu

‡jknowles@econ.sas.upenn.edu.

can verify some simple restrictions on the joint distribution of criminality by race. We also extend these results to the case where the police minimize the rate of unpunished crime.

1 Introduction

Tests for racial discrimination in the US courts tend to ignore the potential role of behavioral models as a tool for identifying bias against minorities. In the case of discretionary searches of motor vehicles, the typical procedure is to attempt to compute search rates by race of the motorist, and then argue from search-rate disparities to racial bias. Two fundamental difficulties in this approach are measuring the composition of the population at risk of search and distinguishing the effects of race from those of other variables correlated with race. These issues are sometimes referred to as 'the denominator problem' and the 'omitted-variables problem', respectively.

A recent paper by Knowles, Persico, and Todd (KPT, 2001) circumvents these difficulties by shifting attention from the analysis of search rates to that of the success rates of search, in the spirit of Becker (1957). Rather than measuring the size of policing disparities, KPT propose a test that distinguishes disparities due to racial bias from those that can be explained by the maximization of the success rate of searches (also known as 'hit rates' or 'find rates') via statistical discrimination. If valid, this test is vastly more convenient than the standard approach of comparing search rates, because it does not require knowledge of the underlying population composition, and it is not invalidated by the possible correlation of race with other predictors of criminality.

Empirically, the KPT model with unbiased policing predicts police will search high-crime-propensity groups with higher probability, and that hit rates will be equated across all observed attributes. Several empirical analyses in addition to KPT have found hit rate differentials by race that are not different from zero at conventional levels of statistical significance.[1] Where hit rates differ, and especially when the differences are statistically significant, they are almost always lower among the groups that appear to be searched with higher probability–that is, among African-American and Hispanic motorists. Examples from the past few years include Missouri and Florida, where state-wide hit rates on Hispanic motorists were about half that of white motorists; Minneapolis, where these searches were just over one-third as successful; and Pennsylvania, where these searches were less than two-thirds as successful. Less extreme, but still substantial, disparities are found when comparing searches of African American motorists to those of white motorists in these jurisdictions.[2] Under the KPT model, the interpretation of such hit-rate disparities is very clear: the police appear to be biased against African-American and Hispanic motorists.

A fundamental issue in the KPT approach is the motivation of unbi-

[1] Some of these studies are summarized in the web page of Lamberth Consulting, under the heading "racial profiling doesn't work" (www.lamberthconsulting.com.).

[2] The Missouri report for 2001, available at the web page of the Attorney-General's office, reveals hit rates of 22% for whites, 15% for blacks, and 11% for Hispanics. In Florida, summarized in Anwar and Fang (2004), the corresponding hit rates are 25%, 21%, and 12%. For Minneapolis, hits rates among discretionary searches are reported to be 13%, 11%, and 5%, respectively. For Pennsylvania, the figures are 29%, 21%, and 17%, as reported in Engle et al. (2004).

ased police. Is their objective really the maximization of the success rate of searches? It is this assumption that generates equality of hit rates in their model, and hence the validity of the test would seem to be linked to the plausibility of this hit-rate maximization (HRM) hypothesis. It is easy to sympathize with those who do not readily accept this assumption, such as Harcourt (2004) and Manski (2005), as the minimization of crime would seem to a more desirable goal, from the point of view of social welfare, than the maximization of arrests. It is would be reassuring if the optimal behavior of police agents motivated by crime-rate minimization (CRM) was approximately the same as under CRM, but in fact Persico (2002) shows that search policies that are optimal under HRM can actually lead to maximization of crime.

One might be tempted to turn to the extensive literature on policing for help. Unfortunately the nature of the police objective function does not seem to play an important role in empirical studies of police behavior. A recent report by the National Research Council (Skogan and Frydl, 2004), despite devoting two large chapters to the determinants of police behavior, never explicitly considers the goals of the police officers themselves. While it seems to be generally agreed that crime minimization should be one of the main goals of the police (see for instance, diIulio, 1993), this does not rule out the possibility that the police actually act so as to maximize the detection and punishment of criminals.

An emphasis on arrest rates might arise for instrumental reasons, because arrest rates are seen to be the best available measure of the police agency's ability to reduce crime, or because arrest rates are the best measure of the crime-reduction effort of the individual cop. But there also seems to be some intrinsic value attached to arrests and convictions. For instance, Alpert and Mclaren (1993) list crime rates as the first of four "enshrined" measures of police performance; however the second measure is arrest rates. They argue: "For most practical purposes, these are the statistics by which police departments throughout the United States are now held accountable."[3] To summarize, there seems to be little evidence to decide between HRM and CRM as a model of police behavior, but a general feeling that the CRM objective better reflects the normative purpose of the police. In testing for biased policing therefore, it would seem prudent to allow for the possibility that police are in fact crime-rate minimizers.

In this paper, we derive conditions that permit an inference of racial bias (i.e., rule out unbiased policing) from observed hit rates when the true objective of police behavior is to minimize crime rather than to maximize arrests. We focus on the simple case with two observable groups. The members of these groups each choose to commit crime if the group-specific search rate falls short of an individual-specific threshold value that reflects the individual's net gains from crime. We ask whether there exist restrictions on the distributions of this crime propensity that permit inference of racial bias from hit-rate data under the CRM hypothesis. In particular, we seek to distinguish between two explanations of search disparities: (1) minorities are more likely to be searched because police are biased and (2) minorities are more

[3] The acting director of the Bureau of Justice Statistics in 1993 cites as evidence of the success of the "War on Drugs" campaign the fact that "drug offenses now account for a larger share of convictions and imprisonments than ever before". Lawrence A. Greenfeld, in the foreword to *Performance Measures for the Criminal Justice System*.

likely to be searched because a greater fraction of them are guilty.

We begin by deriving conditions for an inference of bias were search rates, hit rates, and marginal deterrence effects to be observed. We then present a negative finding on identification for the realistic case when marginal deterrence effects are not observed. In our model, it is the density of the crime-propensity distribution that determines deterrence effects, while the cumulative distribution function determines hit rates. The main identification problem under CRM arises because the optimal strategy of unbiased police is to choose group-specific search rates such that the densities are equated at each group's search rate. In the absence of further assumptions, equating the densities does not restrict the ranking of the CDFs evaluated at these search rates. That is, unlike HRM, unbiased CRM policing does not yield restrictions on hit rate disparities.

We then derive the necessary condition that must be satisfied in our model for an inference of bias to be possible. We show that this condition is satisfied under some simple restrictions on the relationship between the two distributions: if there is a high-crime group whose distribution can be written as a spreading out of that of the low-crime group and if the density of the high-crime group is declining, unbiased police will not choose search rates such that the success rate from searching the high-crime group is less than that of searching the low-crime group. Hence a pattern of higher search rates and lower hit rates for a given group cannot be justified by simply alleging a higher crime propensity of that group.

It may be that what the police actually care about is not crime per se, but rather the rate at which crimes go unpunished. We therefore consider a model in which police attempt to minimize the rate at which people commit crimes without being caught. We refer to this objective as unpunished crime minimization (UCM). Manski (2005) makes a similar distinction between ex ante and ex post search. We find that the same conditions that permit an inference of bias from search rates and hit rates under CRM also permit this inference under UCM. Theoretically, we show that inference under UCM is valid for weaker conditions than our conditions for CRM, but in practice these conditions are likely to be equivalent.

Because we abstract from non-racial observables, we do not address the problem of correlation between other crime-related observables and race. Suppose that, in the real world, non-racial variables were actually not informative about crime probabilities. How useful would our results be? The answer depends on who bears the burden of proof in a case where the police are accused of bias. The spreading condition is equivalent in our model to the condition that the deterrent effect of search on the high-crime group is less responsive than that of the other group. It is not plausible to assume that one can a priori rule out local violations of this restriction over the range of feasible search rates. However, to the extent that the plaintiff is simply required to show that our restrictions are satisfied at observed search rates, and that there is no reason to believe them violated elsewhere, then our results could be used to make the case that the police are biased. Conversely, an effective defense against such a charge would be to show that the high-crime group is in fact more responsive to deterrence than is the low-crime group.

Section 2 presents a simple CRM model of police behavior, taking as given the distributions of crime propensities and, hence, crime-rate response functions. We derive conditions under which the combination of higher search rates and lower hit rates would lead to an inference of biased policing under the CRM hypothesis, and we discuss practical issues for the application of this result. In Section 3, we derive analogous results under the UCM hypothesis. We conclude with some ideas for future research.

2 A CRM Model of Policing

In this section we develop a model of discretionary motor-vehicle searches motivated by crime-rate minimization. We consider a simple world in which there exist one cop and many motorists of different observable types and unobservable crime propensities. The cop is aware only of the motorist type and the distribution of crime propensities within each type; the motorist is aware of the probability of being searched. We ask how observed patterns of search rates and hit rates would depend on whether the cop is biased against some observable group. In other words, is it possible, in our basic model, to infer racial bias from search rate and hit rate statistics, or can the variation in these statistics always be explained by hypothesized variation in crime propensities across groups? We show that, under CRM, the ability to infer bias from police search data requires additional restrictions on the distribution of unobservable variables. We then derive a set of conditions that permit such inference.

2.1 The basic model

For simplicity, suppose there are two observable types of motorists.[4] Let μ_r denote the share of group r in the population, where $r \in \{a, b\}$.

Motorists also differ according to an unobservable characteristic ν which determines whether or not they commit a crime for any given probability of being caught by the cop. Motorist i of group r commits a crime whenever

$$\nu_i > p_r^s$$

where p_r^s denotes the group-r probability of being apprehended if guilty of committing a crime.[5] Note that, for any value $\nu_i < 0$, individual i will not commit crime at any search rate, whereas he will invariably commit crime if $\nu_i > 1$. Suppose the individual crime propensities ν_i in group r are independent draws from a distribution with density ϕ_r and cumulative distribution function Φ_r. We can write the crime probability of the group as a function $f_r(p_r^s) = [1 - \Phi_r(p_r^s)]$. We call this function the *crime-rate*

[4] The extension of the main results to a model with multiple types is quite straightforward. As should be clear from inspection of the forthcoming optimality conditions, these results hold for pairwise comparisons (across groups) of search rates and hit rates. However, as discussed in the Applications section, the extension to mulitple types when the researcher only observes aggregations of these types (i.e., the researcher observes only a subset of the attributes observed by the cop) is not as straightforward and is beyond the scope of this analysis.

[5] In this model, the group-specific probability of being apprehended is just the group-specific probability of search. This result arises because group identity r is the only observed attribute that may determine police search decisions. The result is more general than it may appear, given that group identity may be determined by numerous observed attributes in a model with many groups.

response function of group r. We shall focus attention on the group-specific density ϕ_r, which, by construction, determines the marginal deterrence effect of searches. An important assumption that we maintain throughout is that this density is continuously differentiable.

The number of crimes committed by each group is therefore determined by the search probability and the mass of the group.

$$n_r(p_r^s) = f_r(p_r^s)\mu_r \tag{1}$$

The cop apprehends criminals by searching their motor vehicles for contraband. We assume that all guilty motorists are caught if searched; hence the probability of search is identical with p_r^s as defined above. We assume that the cop can choose to search a motorist with probability zero, and can choose any positive search rate up to a maximum search rate $\overline{p}_r$. This rules out non-discretionary searches and allows for the possibility of non-financial constraints on search, such as the availability or visibility of a given type of motorist. We denote the set of feasible search-rate pairs by P.

We assume the cop is perfectly informed about the crime-rate response function. However, the cop can observe neither whether the individual has committed a crime nor the individual's crime propensity. Hence the cop is only aware of the characteristics of the motorist's group. The success-rate π_r of search on a given group r therefore equals the crime rate within the group:

$$\pi_r\left(p_r^s\right) = f_r(p_r^s)$$

Search is costly; we assume that the cost of search is constant and does not vary by group. We set the search cost equal to one. Hence the budget constraint, given that the cop has resources Y^s, can be written as:

$$\mu_a p_a^s + \mu_b p_b^s = Y^s \tag{2}$$

A general assumption that we maintain throughout is that the budget constraint binds: any search-rate pair (p_a, p_b) that results in zero crime for any group requires resources greater than Y^s.

We assume the cop has preferences over a weighted sum of the number of crimes committed by each group. Let $u(n_a, n_b)$ be a utility function that represents these preferences. We write the cop's problem as choosing group-specific search probabilities p_r^s to minimize u subject to the budget constraint:

$$\min_{p_a^s, p_b^s} u(n_a, n_b) \tag{3}$$

subject to (1) and (2).

We now define the search-rate pair that will be the focus of the analysis and give conditions such that this pair represents an interior solution to the problem (3).

Definition 1 *The feasible set of search rates is denoted by*

$$P = \{(p_a, p_b) \, | p_a \in [0, \overline{p_a}] \text{ and } p_b \in [0, \overline{p_b}] \}$$

, where $\overline{p_a} \in (0,1)$ and $\overline{p_b} \in (0,1)$ represent the highest feasible search rates of each group.

For the purposes of our analysis it is convenient to set the upper bounds $\overline{p}$ equal to 1, but given the concerns raised in the existing literature, such as Dharmapala and Ross (2004), regarding the possibility that search rates are in fact bounded far below one, we introduce this notation so that it is clear that our results do not depend on the assumption that the upper bound is equal to one.

Definition 2 *Let the pair $p^* = (p_a^*, p_b^*)$ denote the search rates that are observed. We assume that $p_a^* > p_b^*$ and that p^* is on the interior of the feasible set P.*

To represent the two main behavioral hypotheses that explain unequal search rates, we now formalise them in the model. Without loss of generality, we suppose, as in the preceding definition, that group a corresponds to the group that is searched at a higher rate in the real world.

The first hypothesis is that the cop is biased against this group. To represent biased policing in our model, we assume that the cop's utility function weights each crime according to the group-membership of the person that has committed the crime. Let the bias parameter $\xi \geq 1$ denote the weight that u places on crime committed by group a, where the weight on crime committed by group b is normalized to 1.

$$u(n_a, n_b) = \xi n_a + n_b \tag{4}$$

An inference of bias against group a, therefore, would arise from a finding of $\xi > 1$.

The competing hypothesis is that group a is a "high-crime" group. This hypothesis requires that, were both groups to be searched at an identical rate, then the crime rate, and hence the success rate of searches, would be higher for group a. Notice that this hypothesis does not restrict the crime rates at the *observed* search-rate pair p^*. Given our model of the crime decision, the high-crime hypothesis requires the assumption that, at any given search rate, the crime-rate probability of group a would be higher than that of group b. This implies the criminality distributions Φ_r can be ranked in the sense of first-order stochastic dominance. For the derivation of our main results, we may restrict attention to these distributions on $[p_b^*, p_a^*]$.

Assumption 1 *Suppose that the distributions satisfy the following restriction on the quantile-quantile plot of Φ_a versus Φ_b, denoted $g(.)$*

$$g(t) \geq t, t \in [p_b^*, p_a^*] \tag{5}$$

where

$$\Phi_b(t) = \Phi_a(g(t))$$

That is, Φ_a first-order stochastically dominates Φ_b on $[p_b^*, p_a^*]$.

In the analysis to follow, the goal is to identify conditions that permit the observer to distinguish between these two explanations of search disparities on the basis of search rates and success rates of searches. We take it as given that the high-crime group is known to be the one that is searched at higher rate, and ask whether we can rule out unbiased policing on the basis of observed hit rates. That is, assuming that the high-crime hypothesis is correct, is the data generating process sufficient to permit identification of a lower bound on ξ that exceeds 1? If not, than what additional information would suffice to permit such identification and, therefore, potentially lead to a valid inference of bias based on observed variation in search rates and hit rates?

2.2 Inference from Marginal Deterrence Effects

Suppose we observe that the cop is searching both groups with interior probabilities. Given this observation, how can we infer anything about racial bias, according to our model? The discussion proceeds by first establishing necessary conditions such that a pair of interior probabilities p^* may indeed be an optimum, and showing that any deviation from equality of marginal deterrence effects at p^* must be due to bias. This requires only local regularity assumptions. Then we explore the conditions needed to permit inference about bias from observable statistics, such as search rates and success rates of search.

Given the assumption that p^* is interior, we mainly need to ensure that the crime-rate response functions are continuously differentiable at p^*.

Assumption 2 *For* $r \in \{a, b\}$, $\phi_r(p_r^*)$ *is continuous and strictly positive.*

Proposition 1 *Suppose that Assumption* 2 *is satisfied. Then the interior search-rate pair* (p_a^*, p_b^*) *is a local optimum only if it satisfies the first-order condition of* (3)*:*

$$\frac{f_a'(p_a^*)}{f_b'(p_b^*)} = \frac{1}{\xi} \tag{6}$$

Notice that we have not established whether p^* minimizes the weighted crime rate, because we have not ruled out other local optima. Our immediate concern is rather, given that we observe p^*, we could infer racial bias against group a if the marginal deterrence effects $f_r'(p_r^*)$ were observed to satisfy the following inequality:

$$\frac{f_a'(p_a^*)}{f_b'(p_b^*)} < 1 \tag{7}$$

Corollary 2 *Suppose that Assumption* 2 *is satisfied. Then, at the interior search-rate pair* (p_a^*, p_b^*) , *inequality* (7) *cannot arise from unbiased police who minimize crime, as defined by* (3) .

Even if there were other local optima that dominated p^* for unbiased cops, condition (7) would indicate, under our model, bias against group a. This result is clear from Proposition 1, where unbiased policing has $\xi = 1$ and, therefore, $f_a'(p_a^*) = f_b'(p_b^*)$ at any interior solution. Hence, inference of bias from marginal deterrence effects requires only the assumption that the densities are smooth and positive in some neighborhood of p^*. However we could not infer non-bias against a from violation of (7) , because the possibility of other optima preferred by unbiased police cannot be ruled out. These results form the foundation for our assessment of conditions under which we can infer bias from police search data.

2.3 Inference from Hit Rates

The results in the previous section cannot be directly applied to the real-life problem of testing for bias, because it is extremely unlikely that it will be possible to observe the marginal deterrence effects. Suppose instead that we observe only the search rates and the success rates of those searches, for each observable group of motorist. What additional conditions are required to make inferences about bias from those statistics, under the above assumptions?

It is clear that inference from observable statistics is only possible under CRM to the extent that it is possible to infer something about the densities of ν from these statistics. Without further restrictions on the distributions however, the hit rates only tell us the area under ϕ_r to the right of the observed search probability p_r^*, not the height of ϕ_r. This negative finding may be stated as follows:

Proposition 3 *Assumptions* 1 *and* 2 *are not sufficient to allow inference of racial bias from observation of search rates* (p_a^*, p_b^*) *and hit rates* $(\pi_a(p_a^*), \pi_b(p_b^*))$, *when police are known to minimize crime, as defined by* (3).

2.4 Restrictions on Distributions

We have established conditions on marginal deterrence effects that an interior solution must satisfy if the cop is unbiased. The purpose of this section is to establish weak mathematical conditions that make inference of bias possible from observables. We assume that the data reveal which group is experiencing the higher search rate, and which is experiencing the higher hit rate. We assume that, as is likely in practice, the analyst is aware of neither the functional forms of the crime-propensity distributions nor the absolute level of the search rates. While it may be reasonable to assume that the hit rates are precisely measured, it should be clear that this additional information is of little use in the context of the CRM model, given the above limitations on the analyst's knowledge.

The observation that group a is being searched at a higher rate than group b and yet is yielding a lower success rate of searches is represented in our model by $p_a^* > p_b^*$ and $\pi_a(p_a^*) < \pi_b(p_b^*)$. As implied by Proposition 3, an inference of bias from such an observation will require further assumptions. We show here that joint restrictions on the quantile-quantile plot g and the densities ϕ_r will be necessary for such inference.

Proposition 4 *It is possible to infer bias against group a from knowledge that $p_a^* > p_b^*$ and $\pi_a(p_a^*) < \pi_b(p_b^*)$, given Assumptions 1 and 2, under crime minimization (3), if and only if the following proposition is true:*[6]

$$p_a^* > g(p_b^*) \Rightarrow g'(p_b^*) > \frac{\phi_a(p_a^*)}{\phi_a(g(p_b^*))}$$

Note that if $p_a^* > g(p_b^*)$ implies $\phi_a(p_a^*) < \phi_a(g(p_b^*))$ then this proposition implies that inference is only possible if g' is bounded below by 1. This is because we cannot rule out p_a^* and $g(p_b^*)$ being arbitrarily close, and hence the ratio of the densities cannot, *a priori,* be restricted to be less than one. In the corollary, we formalize this as a pair of restrictions that jointly suffice for inference of bias when the high-crime group is searched at a higher rate and these searches are less successful, without requiring knowledge of the specific functions forms of ϕ_r and g:

Corollary 5 *Suppose that Assumptions 1 and 2 are satisfied. Then, at interior search rates p^*, knowledge of $g'(p_b^*) \geq 1$ and $\phi_a(p_a^*) < \phi_a(g(p_b^*))$ ensures that higher search rates $p_a^* > p_b^*$ and lower hit rates $\pi_a(p_a^*) < \pi_b(p_b^*)$ cannot arise from unbiased police who minimize crime under CRM.*

The proposition and its corollary require local restrictions on g and ϕ_a. However, the difficulty of observing actual search rates leads us to believe that any attempt to verify these restrictions would require that they hold over a range of feasible search rates. Moreover, when such a requirement is

[6] As noted in the proof contained in the Appendix, the combination of stochastic dominance (Assumption 1) and lower hit rates implies higher search rates of the high-crime group $p_a^* > g(p_b^*) \geq p_b^*$. In contrast, stochastic dominance and higher search rates imply no restrictions on relative hit rates.

made, the substance of the restrictions becomes clearer. That is, inference of bias from unequal hit rates is valid when the high-crime distribution is known to (i) be more spread out to the right than the distribution of the low-crime group and (ii) exhibit a declining density over $[g(p_b^*), p_a^*]$.[7]

2.4.1 Examples of Restrictions on Distributions

Our analysis proceeds by briefly illustrating what these restrictions mean for the shapes of the crime propensity distributions and for the relationship between these distributions. We focus on two cases. First, we consider a *shift and spread*, which means that $g'(t) \geq 1$ for all $t \in [0,1]$. As the above analysis indicates, knowledge of $g'(p_b^*) \geq 1$ may enable inference of bias from observed hit rates. As search rates are difficult to measure, empirical verification of this condition may require that the inequality hold over the range of plausible search rates, so we use the entire unit interval for illustration. Second, to exemplify a relationship among distributions that would not enable such an inference, we consider what we call a *shift and tightening*, in which case $0 < g'(t) < 1$ for all $t \in [0,1]$. One may interpret the shift-and-spread restriction as a case in which members of low-crime group are more similar in terms of crime propensity than are members of high-crime group. The converse is true under the shift-and-tightening restriction. In both cases we maintain the assumptions that the group-a distribution stochastically dominates the group-b distribution and the group-a density is declining on $[0,1]$.

In both cases, we restrict attention to distributions that are necessarily

[7]Note that stochastic dominance and lower hit rates imply $g(p_b^*) < p_a^*$.

declining over $[p_b^*, p_a^*]$, as occurs, for example, whenever ϕ_a is unimodal with mode less than p_b^*. Should this density be allowed to increase above p_b^* or, in particular, above $g(p_b^*)$, then inference is not possible without more information on the distributions and, perhaps, search rates. As should be clear from the derivations above, an interior solution for unbiased policing satisfies

$$g'(p_b^*) = \frac{\phi_a(p_a^*)}{\phi_a(g(p_b^*))}$$

If ϕ_a has a mode or modes above $g(p_b^*)$, then, without additional information, the fraction on the right cannot be bounded below unity based on observed search rates and hit rates.

Restriction 1: Shift and Spread Suppose that

$$g(t) \geq t \text{ and } g'(t) \geq 1 \text{ for all } t \in [0,1] \tag{8}$$

Here, the group-a distribution is both shifted up the real line and more widely spread. Thus, the group-a crime-rate function $f_a(p_a^s)$ is less responsive to changes in the search rate than is the group-b crime-rate function. Graphically, the downward sloping portion of the group-a density may lie everywhere above the group-b density (as in Figure 1a) or may perhaps start below the group-b density and cross it on $[0,1]$ (Figure 1b). This relationship among crime-propensity distributions necessarily satisfies $g'(p_b^*) \geq 1$ and yields the following result:

Proposition 6 *Suppose that group a has a distribution of crime propensities ν that is a shift and spread of that of group b, such that (8) holds. Suppose further that Assumption 2 is satisfied and ϕ_a is strictly decreasing on $[0,1]$.*

Then, unbiased police who minimize crime, as defined by (3), *choose* $p_a^* \leq g(p_b^*)$. *Therefore, the hit rate among searches of the high-crime group is weakly greater than the hit rate for the low-crime group, with strict inequality unless* $g'(p_b^*) = 1$. *A lower hit rate only arises if police are biased against group* a.

Restriction 2: Shift and Tightening Now consider the case where

$$g(t) \geq t \text{ and } g'(t) < 1 \text{ for all } t \in [0, 1] \tag{9}$$

In this case, the group-a crime propensity distribution is tighter over $[0, 1]$ and the crime-rate function is therefore more responsive to changes in the search rate than is the group-b crime rate. In Figure 2, note that the group-a density lies above the group-b density but declines more quickly on $[0, 1]$. This relationship among crime-propensity distributions obviously violates $g'(t) \geq 1$ for any search rate p_b^*. With these conditions, we instead have the following result on hit rates:

Proposition 7 *Suppose that the distribution of the high-crime group is a shift and tightening of that of the low-crime group, such that* (9) *holds. Suppose further that Assumption* 2 *is satisfied and* ϕ_a *is strictly decreasing on* $[0, 1]$. *Then, unbiased police who minimize crime, as defined by* (3), *choose* $p_a^* > g(p_b^*) > p_b^*$. *Therefore, the high-crime group is searched at a higher rate, and the hit rate for this group is lower than for the low-crime group.*

Under the conditions required for Proposition 6, it is possible to infer that police are biased against group a. If these conditions are violated, as in Restriction 2, then we cannot rule out unbiased crime minimization as an explanation of the higher search rates of group a. However, under Restriction 2, inspection of Proposition 7 reveals that observing higher find rates for the high-crime group would indicate a bias against group b. By restricting $\xi \geq 1$ in our model, we have not allowed for this possibility, so that we may focus attention on what appears to be the prevalent concern in analyses of police search data.

2.5 Applications

The findings thus far draw attention to the difficulty of drawing an inference of biased policing based on the observation that search rates are higher and hit rates are lower among racial and ethnic minority drivers than among white drivers. We have established conditions under which such an inference may be made even if the minority group is said to constitute a high-crime-propensity group. Empirical verification of these conditions, however, poses a serious challenge. In contrast, application of the KPT hit-rate test appears straightforward, should one accept the HRM hypothesis. It is worth noting that the relationship between hit-rate outcomes and bias under HRM and is observationally equivalent to that under CRM if the crime-propensity distributions are identical, but for a shift in location–that is, $g'(t) = 1$ for all $t \in [0, 1]$.[8] Of course, this assumption is very strong and unlikely to be generally acceptable.

It may be more acceptable to assume that the distribution of the high-crime group is a shift and spread, in the sense defined above. An interpre-

[8]Proof of this result follows the same line of reasoning as that applied in the proof of Proposition 6.

tation of the condition $g'(.) \geq 1$ is that members of the high-crime group are less similar to each other in terms of the net returns to crime than is the case for members of the low-crime group. This is a plausible condition that could in principle be empirically verified.[9] Inference in this case requires very general assumptions about the functional form of the high-crime group distribution: positive but decreasing density. However, the difficulty of precisely measuring search rates–e.g., the denominator problem–suggests to us that the restrictions on $g'(.)$ and ϕ_a need be verified to hold for all plausible search rates. Further, to be literally implemented, the test requires measuring not the distributions of criminal propensities per se, but rather the perceptions of the cop of these distributions over plausible search rates.

Another important difference in terms of observable results between the CRM model and the KPT test is that, even under conditions that would permit testing conditional on observables, the CRM model does not justify integration over variables that are observed by the cops but not by the analyst. In contrast, the KPT equilibrium with unbiased policing results in all motorists having the same guilt rate so that aggregation is possible. Under the CRM equilibrium, such aggregation would require additional restrictions on the distribution of these unobserved (by the analyst) attributes and/or within-group variation in crime-propensity distributions conditional on these attributes. Thus, the omitted variable problem faced by researchers who assess search rate disparities arises here as well.

[9]The empirical analysis in Persico (2002) may be interpreted as such, taking the reported income distribution as a proxy for the distribution of propensity to not commit crime. In Persico's model, individuals differ only in legitimate earnings opportunities, so this interpretation would be appropriate. His results suggest, in fact, that $g'(t) < 1$.

Persico (2002) analyzes fairness and effectiveness of policing, maintaining as we do the information structure of KPT and the focus on interior solutions. Persico posits variation in legal earnings opportunities available to members of each group as the sole source of unobserved heterogeneity in his model, and hence the sole source of variation in propensity to (not) commit crime. He then asks whether a trade-off exists between fairness and efficiency, which he implements as a question of whether a marginal decrease in search rate disparities from the no-bias KPT equilibrium would increase aggregate crime. He derives a key restriction on the quantile-quantile plot between the group-specific distributions of earnings, which is analogous to our shift-and-tightening restriction on crime propensity distributions.[10] His empirical analysis of reported earnings suggests that this restriction holds.

So far we have only considered one alternative to HRM, the CRM hypothesis. While we find this a plausible and beneficial motivation for the cop, one might be concerned more about minimizing the rate of crime that goes unpunished, rather than just the crime rate. Under CRM, catching criminals is only instrumental in reducing the crime rate, not desirable in itself. In the next section, we extend the model to allow for the apprehension of criminals. This is a natural extension, because there is a sense in which it combines both the HRM and the CRM models of police behavior.

[10] This result may be extended to cover large changes in search rates, if the densities are both known to be declining over a large range of search rates. See Appendix.

3 The UCM Model of Policing

What if cops seek not to just either deter crime or to apprehend criminals, as in CRM or HRM, respectively, but instead seek to minimize the rate at which people commit crimes without being caught? This type of motivation could arise because catching criminal signals higher diligence or competence of the cop making the arrest, or because catching criminals has some intrinsic value to the community. We now extend the model by modifying the cop's objective function to combine both the deterrence and apprehension objectives. We refer to this model as unpunished crime minimization (UCM).

We write the cop's problem as choosing group-specific search probabilities p_r^s to minimize the weighted sum of unpunished crimes:

$$\min_{p_a^s, p_b^s} \ \xi \left[n_a(p_a^s) \cdot (1 - p_a^s) \right] + n_b(p_b^s) \cdot (1 - p_b^s) \tag{10}$$

This simple modification of the CRM objective function (4) arises from the independence of guilt and search probabilities, conditional on group.[11]

Suppose again that we could observe marginal deterrence effects, in addition to search rates and hit rates, and wish to make inferences on racial bias.

Proposition 8 *Suppose that Assumption* 2 *is satisfied. Then the interior search-rate pair* (p_a^*, p_b^*) *is a local optimum only if it satisfies the first-order*

[11] As in out analysis of CRM, group summarizes all information available to the cops about individual crime propensity. Therefore, conditional on group r, the rate of search among the guilty is just the marginal search rate p_r^s, so the fraction of group-r criminals who are not apprehended is just $(1 - p_r^s)$.

condition of (10)*:*

$$\frac{(1-p_a^*)\, f_a'(p_a^*) - f_a(p_a^*)}{(1-p_b^*)\, f_b'(p_b^*) - f_b(p_b^*)} = \frac{1}{\xi} \tag{11}$$

or, equivalently,

$$\frac{(1-p_a^*)\, \phi_a(p_a^*) + \pi_a(p_a^*)}{(1-p_b^*)\, \phi_b(p_b^*) + \pi_b(p_b^*)} = \frac{1}{\xi}$$

Notice that the first component of the numerator and denominator of (11) includes the main element of the CRM optimality condition–the marginal deterrent effect–and the second component includes the main element of the HRM optimality condition–the hit rate. We may see this optimality condition as a combination of the optimality conditions under CRM and HRM.

What can we infer from search rates and hit rates under UCM? Recognizing again that observation of marginal deterrent effects is highly implausible, we seek to determine conditions under which we could infer bias against group a based on observation that this group is searched at a higher rate and the success rate of these searches is lower.

To do so, we consider the optimality condition for the unbiased cop ($\xi = 1$) who minimizes unpunished crime, as in (10). We may rewrite (11) to identify the following condition for an interior solution (p_a^*, p_b^*) for unbiased policing:

$$\pi_a(p_a^*) - \pi_b(p_b^*) = (1-p_b^*)\, \phi_b(p_b^*) - (1-p_a^*)\, \phi_a(p_a^*) \tag{12}$$

Inspection of (12) yields the following result:

Proposition 9 *Suppose that Assumptions* 1 *and* 2 *are satisfied. Then higher search rates* $p_a^* > p_b^*$ *and weakly lower hit rates* $\pi_a(p_a^*) \leq \pi_b(p_b^*)$ *cannot arise from unbiased police who minimize unpunished crime, as defined by* (10)*, when* $g'(p_b^*) \geq 1$ *and* $\phi_a(p_a^*) < \phi_a(g(p_b^*))$*.*

Thus, once again, higher search rates and lower hit rates among the high-crime group cannot arise in the absence of bias when the high-crime group distribution is a spread of the low-crime group distribution and the density is declining. In fact, as can be derived from the proof in the appendix, the spread restriction may be weakened somewhat. In particular, the main result of Proposition 9 holds for any $g'(p_b^*) > \left(\frac{1-p_a^*}{1-p_b^*}\right)$. This weaker constraint is not particularly applicable, as search rates are notoriously difficult to measure.

It seems likely that police are directly concerned with both crime deterrence and apprehension of criminals (e.g., for the purposes of prevention, incapacitation, and/or retribution), so it is reassuring that key findings under CRM apply under UCM. It should be noted that previous findings on optimal policing in the absence of bias must be reconsidered, if the UCM model is more appropriate. As noted above, Persico (2002) develops conditions–closely related to the shift-and-tightening condition in our framework–under which a marginal increase in fairness will increase the aggregate crime rate, yielding a trade-off between fairness and efficiency. Should efficiency instead be defined with respect to the rate of unpunished crimes rather than all crimes, these conditions would likely need be to revised.[12]

[12] In out results, the key restriction on the quantile-quantile plot for unbiased policing to yield $\pi_a(p_a^*) < \pi_b(p_b^*)$ changes from $g'(p_b^*) < 1$ (under CRM) to $g'(p_b^*) \leq \left(\frac{1-p_a^*}{1-p_b^*}\right) < 1$ (under UCM). Therefore, it seems that Persico's conditions for the existence of a fairness-efficiency tradeoff would become more stringent as well.

4 Conclusions

This paper explored the potential to distinguish, on the basis of search data, between two common explanations of the apparently high search rates of African-American and Hispanic minorities by police. These competing explanations are that the police are biased against minorities, and that minorities contain a higher share of individuals inclined to commit crime. Accordingly, we considered a basic model where there are two groups of motorists, one of which was assumed to have a higher average value of an idiosyncratic propensity to commit crime. The goal of our analysis was to establish conditions on the distribution of this random crime 'propensity' across groups such that the hit-rate test developed by KPT in the context of the HRM model of police behavior would also be valid if police minimize crime (the CRM model).

We found that the hit-rate test applies if the distributions are invariant in shape across groups. However this condition is too restrictive to be very useful. When distributions are allowed to differ more generally, we found a precise mathematical condition that must be satisfied to permit inference of bias against a group that has a lower success rate of search. We then showed that this condition is satisfied when the high-crime group has a more spread-out distribution, so that it is less responsive at the margin to the deterrent effects of search.

We also extended the model to allow for the possibility that what the police care about is the number of crimes that go unpunished. We called this the UCM model. We found that the same conditions that permit inference from hit rates in the CRM model are also sufficient in the UCM model. Thus, the hit-rate methodology appears somewhat robust to specification of the police objective function, provided our restrictions on distributions are satisfied.

Our results do not however extend a carte blanche to the hit-rate methodology, even when these restrictions are satisfied. Our model, unlike that of KPT, does not imply that guilt-rates will be equalized across observable groups. In real life, groups that are identical in the data will differ according to other attributes observable to the police that can help predict crime. If these attributes are not recorded in the data, then further work is needed to characterize the conditions that will permit valid inferences of bias in such situations.

5 Appendix

5.1 Proofs

Proof. *Proposition 3.* For an arbitrary interior search-rate pair p^*, we know the hit rates are $1 - \Phi_a(p_a^*)$ and $1 - \Phi_b(p_b^*)$. We also know $\phi_a(p_a^*) > 0, \phi_b(p_b^*) > 0$. The second-order condition requires that no more than one density be increasing, but no informative restrictions are placed on $\phi_a(p_a^*)$, $\phi_b(p_b^*)$. Thus, it may be the case that $\phi_a(p_a^*) = \phi_b(p_b^*)$, which satisfies the first-order condition for unbiased policing. For example, Figures 1 and 2 depict distributions that can capture all 3 possible combinations of hit rate and search rate rankings based on unbiased policing when group-a is the high-crime group. Only 3 possible combinations exist, because first-order stochastic dominance rules out any possibility that $\pi_a(p_a^*) < \pi_b(p_b^*)$ when $p_a^* < p_b^*$. Defining $p_a^* > p_b^*$, as we do in the text, rules out one more combination, leaving just the two possibilities on which we focus in Figure 1 ($\pi_a(p_a^*) > \pi_b(p_b^*)$) and Figure 2 ($\pi_a(p_a^*) < \pi_b(p_b^*)$). ■

Proof. *Proposition 4.* By Assumption 1, we may write $\Phi_b(t) = \Phi_a(g(t))$, where $g(t) \geq t$ for all $t \in [p_b^*, p_a^*]$. Thus, $\pi_a(p_a^*) < \pi_b(p_b^*)$ implies $\Phi_a(p_a^*) > \Phi_a(g(p_b^*))$. By monotonicity of the CDF, this inequality requires

$$p_a^* > g(p_b^*)$$

and stochastic dominance requires

$$g(p_b^*) \geq p_b^*$$

yielding the ordering

$$p_a^* > g(p_b^*) \geq p_b^*$$

Next, by taking derivatives of $\Phi_b(t)$ and $\Phi_a(g(t))$ with respect to t, we know:

$$\phi_b(t) = \phi_a(g(t)) \times g'(t)$$

In Corollary 2, under Assumption 2, we have established condition (7) for an inference of racial bias against group a. Condition (7) may be rewritten here as

$$\frac{\phi_a(p_a^*)}{\phi_b(p_b^*)} < 1$$

or, by plugging in for $\phi_b(p_b^*)$,

$$\frac{\phi_a(p_a^*)}{\phi_a(g(p_b^*)) \times g'(p_b^*)} < 1$$

Thus, to infer biased policing from $p_a^* > p_b^*$ and $\Phi_a(p_a^*) > \Phi_b(p_b^*)$, we require that $p_a^* > g(p_b^*)\,(\geq p_b)$ implies

$$\phi_a(p_a^*) < \phi_a(g(p_b^*)) \times g'(p_b^*)$$

If this implication does not arise, then condition (7) need not hold. ■

Proof. *Proposition 6.* An interior solution with no bias ($\xi = 1$) is characterized by

$$\begin{aligned}\phi_a(p_a^*) &= \phi_b(p_b^*) \\ &= \phi_a(g(p_b^*))\, g'(p_b^*)\end{aligned}$$

With $g'(t) \geq 1$ for all $t \in [0,1]$, we know

$$\phi_a(p_a^*) \geq \phi_a(g(p_b^*))$$

with strict inequality unless $g'(p_b^*) = 1$. Given a declining density ϕ_a, we can deduce that $p_a^* \leq g(p_b^*)$. Thus, we know

$$\Phi_a(p_a^*) \leq \Phi_a(g(p_b^*)) = \Phi_b(p_b^*)$$

and

$$\pi_a(p_a^*) \geq \pi_b(p_b^*)$$

with strict inequality unless $g'(p_b^*) = 1$. If instead the cop is biased ($\xi > 1$), then

$$\phi_a(p_a^*) < \phi_a(g(p_b^*))\, g'(p_b^*)$$

and, perhaps

$$\pi_a(p_a^*) < \pi_b(p_b^*)$$

when $g'(t) \geq 1$ for all $t \in [0,1]$ ■

Proof. *Proposition 7.* An interior solution with no bias ($\xi = 1$) is characterized by

$$\phi_a(p_a^*) = \phi_a(g(p_b^*))\, g'(p_b^*)$$

With $g'(t) < 1$, for all $t \in [0,1]$ we know

$$\phi_a(p_a^*) < \phi_a(g(p_b^*))$$

We can deduce that

$$p_a^* > g(p_b^*) > p_b^*$$

That is, the group-a search rate exceeds the group-b search rate. Further,

$$\Phi_a(p_a^*) > \Phi_a(g(p_b^*)) = \Phi_b(p_b^*)$$

Thus,

$$\pi_a(p_a^*) < \pi_b(p_b^*)$$

■

Proof. *Proposition 9.* Suppose hit rates are weakly lower in group a:

$$\pi_a(p_a^*) \leq \pi_b(p_b^*)$$

Therefore, by (12), an interior solution without bias requires

$$(1 - p_a^*)\,\phi_a(p_a^*) \geq (1 - p_b^*)\,\phi_b(p_b^*)$$

Plug in $\phi_a(g(p_b^*))\,g'(p_b^*)$ for $\phi_b(p_b^*)$ to get

$$\phi_a(p_a^*) \geq \frac{(1 - p_b^*)}{(1 - p_a^*)}\phi_a(g(p_b^*))\,g'(p_b^*)$$

Note that, with $p_a^* > p_b^*$, this weak inequality requires

$$\phi_a(p_a^*) > \phi_a(g(p_b^*))\,g'(p_b^*)$$

We impose the restriction that $g'(p_b^*) \geq 1$, so we now require

$$\phi_a(p_a^*) > \phi_a(g(p_b^*))$$

But this condition contradicts the stated restriction on the density. ■

5.2 Do fairness restrictions increase crime?

In order to relate our results to existing literature, it may be helpful to consider a graphical representation in terms of police indifference curves and the feasible set of crime rates. In particular, we introduce what we call the crime possibilities frontier (CPF) and we use this construct to extend previous results on marginal improvements in fairness of search rates. With restrictions on the CPF that arise from assumptions made in the preceding analysis, we are able to easily assess the impact of large changes in relative search rates on the aggregate crime rate.

We know that the no-bias KPT equilibrium has equal hit rates. To generate hit-rate equalization, the high-crime group is searched at a higher rate than the low-crime group. We have found in our analysis that a crime-minimizing interior solution has a high-crime group hit rate greater than (less than) the low-crime group hit rate when the criminality distribution is more (less) diffuse, as summarized by the sign of $(g'(t) - 1)$. We now consider how the sign of $(g'(t) - 1)$ determines the shape of the crime possibilities frontier and the location of the crime-minimizing outcome relative to the KPT equilibrium outcome. We then use this framework to answer the question asked by Persico (2002): Will an increase in fairness (i.e., more equal search rates) from the no-bias KPT equilibrium increase or decrease aggregate crime?

First, note that the budget constraint and the crime-rate response functions define the feasible set of crime rates by group. The boundary of this set, which we call the CPF, has a slope given by $-\frac{\mu_b}{\mu_a}\frac{f'_a}{f'_b}$ or $-\frac{\mu_b}{\mu_a}\frac{\phi_a}{\phi_b}$. The police preferences, as defined by (3), imply linear indifference curves with slope

$-\frac{\mu^b}{\mu^a}\frac{1}{\xi}$, with utility increasing towards the origin. For ease of exposition, suppose that $\mu_a = \mu_b$. Thus, an interior solution to the crime minimization problem occurs where the slope of the CPF is

$$-\frac{1}{\xi}$$

Note how, all else equal, variation in the bias parameter ξ yields rotation of the police indifference curves and movement of the optimum along the frontier. For unbiased cops, the indifference curves have slope -1 and, importantly, every point along the curve yields an identical aggregate crime rate. The closer the indifference curve is to the origin, the lower is the aggregate crime rate. Examples are presented in Figures 3a, 3b, and 3c.

In the absence of bias, the KPT equilibrium is located where the 45-degree line from the origin crosses the frontier. Figure 3c displays a special case in which the KPT equilibrium is the crime-minimizing outcome for unbiased cops. This case arises under a *shifted distributions* restriction, where $g'(.) = 1$.

Suppose we assume that both densities are strictly decreasing for all feasible search rates. Then, we know that the CPF is strictly convex to the origin.[13] Strict convexity of the CPF allows us to easily assess whether more equal search rates increase or decrease aggregate crime relative to the no-bias KPT equilibrium.

Proposition 10 *Suppose the crime possibilities frontier is strictly convex to the origin. (a) Suppose further that the high-crime group crime-propensity*

[13] To prove convexity, note that the slope of the CPF at any point is $-\frac{\mu_b}{\mu_a}\frac{\phi_a}{\phi_b} = -\frac{\phi_a}{\phi_b}$. As f_b increases (Φ_b decreases) along the CPF, f_a decreases (Φ_a increases). We know then, given strictly declining densities, that ϕ_b increases and ϕ_a decreases.

distribution is shifted and spread relative to the low-crime group distribution. Then any increase in fairness from the KPT no-bias equilibrium search rates will decrease aggregate crime, up until the search rates at which aggregate crime is minimized. (b) Suppose instead that the high-crime group crime-propensity distribution is shifted and (weakly) tightened relative to the low-crime group distribution. Then any increase in fairness from the KPT no-bias equilibrium search rates will increase aggregate crime.

Suppose $g'(t) > 1$–Restriction 1, depicted in Figures 1a and 1b. Here, the no-bias ($\xi = 1$) crime-rate minimizing outcome is northwest of the no-bias KPT equilibrium on the CPF. An increase in fairness that decreases group-a search rates will represent a northwest move along the frontier from the KPT equilibrium. This move brings us closer to the crime-rate minimizing outcome. Given the convexity of the frontier, such a move must yield a point on a lower indifference curve with slope -1. Therefore, the aggregate crime rate decreases until one moves past the crime-rate minimizing outcome.

Suppose instead that $0 < g'(t) < 1$–Restriction 2 depicted in Figure 2. Here, the no-bias crime-minimizing outcome is southeast of the no-bias KPT equilibrium on the frontier. Any increase in fairness will again represent a northwest move along the frontier, which is now a move away from the crime-minimizing outcome. Given the convexity of the frontier, such a move must yield a point on a higher indifference curve with slope -1. Therefore, the aggregate crime rate increases.

Hence any policy that causes the cop to equalize search rates across race will increase equilibrium crime if $g' < 1$ and decrease crime otherwise. This finding is an extension of Proposition 2 in Persico (2002), which states that, when F_b is a stretch of F_a, a marginal change from the no-bias KPT equilibrium toward fairness increases crime–i.e., a trade-off between fairness and efficiency exists. Under our Restriction 2, in addition to conditions guaranteeing a strictly convex CPF, this result holds beyond marginal changes.

6 References

References

[1] Alpert, Geoffrey and Mark H. Moore (1993), "Measuring police performance in the New Paradigm of Policing", in *Performance Measures for the Criminal Justice System,* U.S. Department of Justice.

[2] Antonovics, K., and B. Knight (2004) "A New Look at Racial Profiling: Evidence from the Boston Police Department" NBER Working Paper Number.

[3] Anwar and H. Fang (2004). "An Alternative Test of Racial Prejudice in Motor Vehicle Searches: Theory and Evidence" unpublished manuscript.

[4] Becker, Gary S (1957) *The Economics of Discrimination,* University of Chicago Press.

[5] Dharmapala and Ross (2004) "Racial Bias in Motor Vehicle Searches: Additional Theory and Evidence"

[6] DiIulio, John J. (1993), "Rethinking the Criminal Justice System: Towards a New Paradigm", in *Performance Measures for the Criminal Justice System,* U.S. Department of Justice.

[7] Engle, R. et al (2004) *Project on Police-Citizen Contacts: Year 1 Final Report*

[8] Harcourt, B. (2004). "Rethinking Racial Profiling: A Critique of the Economics, Civil Liberties, and Constitutional Literature, and of Criminal Profiling More Generally", *The University of Chicago Law Review*, 71:1275-1381.

[9] Knowles, J., N. Persico, and P. Todd (2001). "Racial Bias in Motor-Vehicle Searches: Theory and Evidence" *Journal of Political Economy*, 109(1), 203-229.

[10] Manski, C. (2005) "Search Profiling with Partial Knowledge of Deterrence," unpublished manuscript.

[11] Persico, N. (2002). "Racial Profiling, Fairness, and Effectiveness of Policing" *American Economic Review* 92(5), 1472-1497

[12] Wesley Skogan and Kathleen Frydl, editors (2004). *Fairness and Effectiveness in Policing: The Evidence,* National Academies Press, Washington DC.

Figure 1a: Shift and Spread

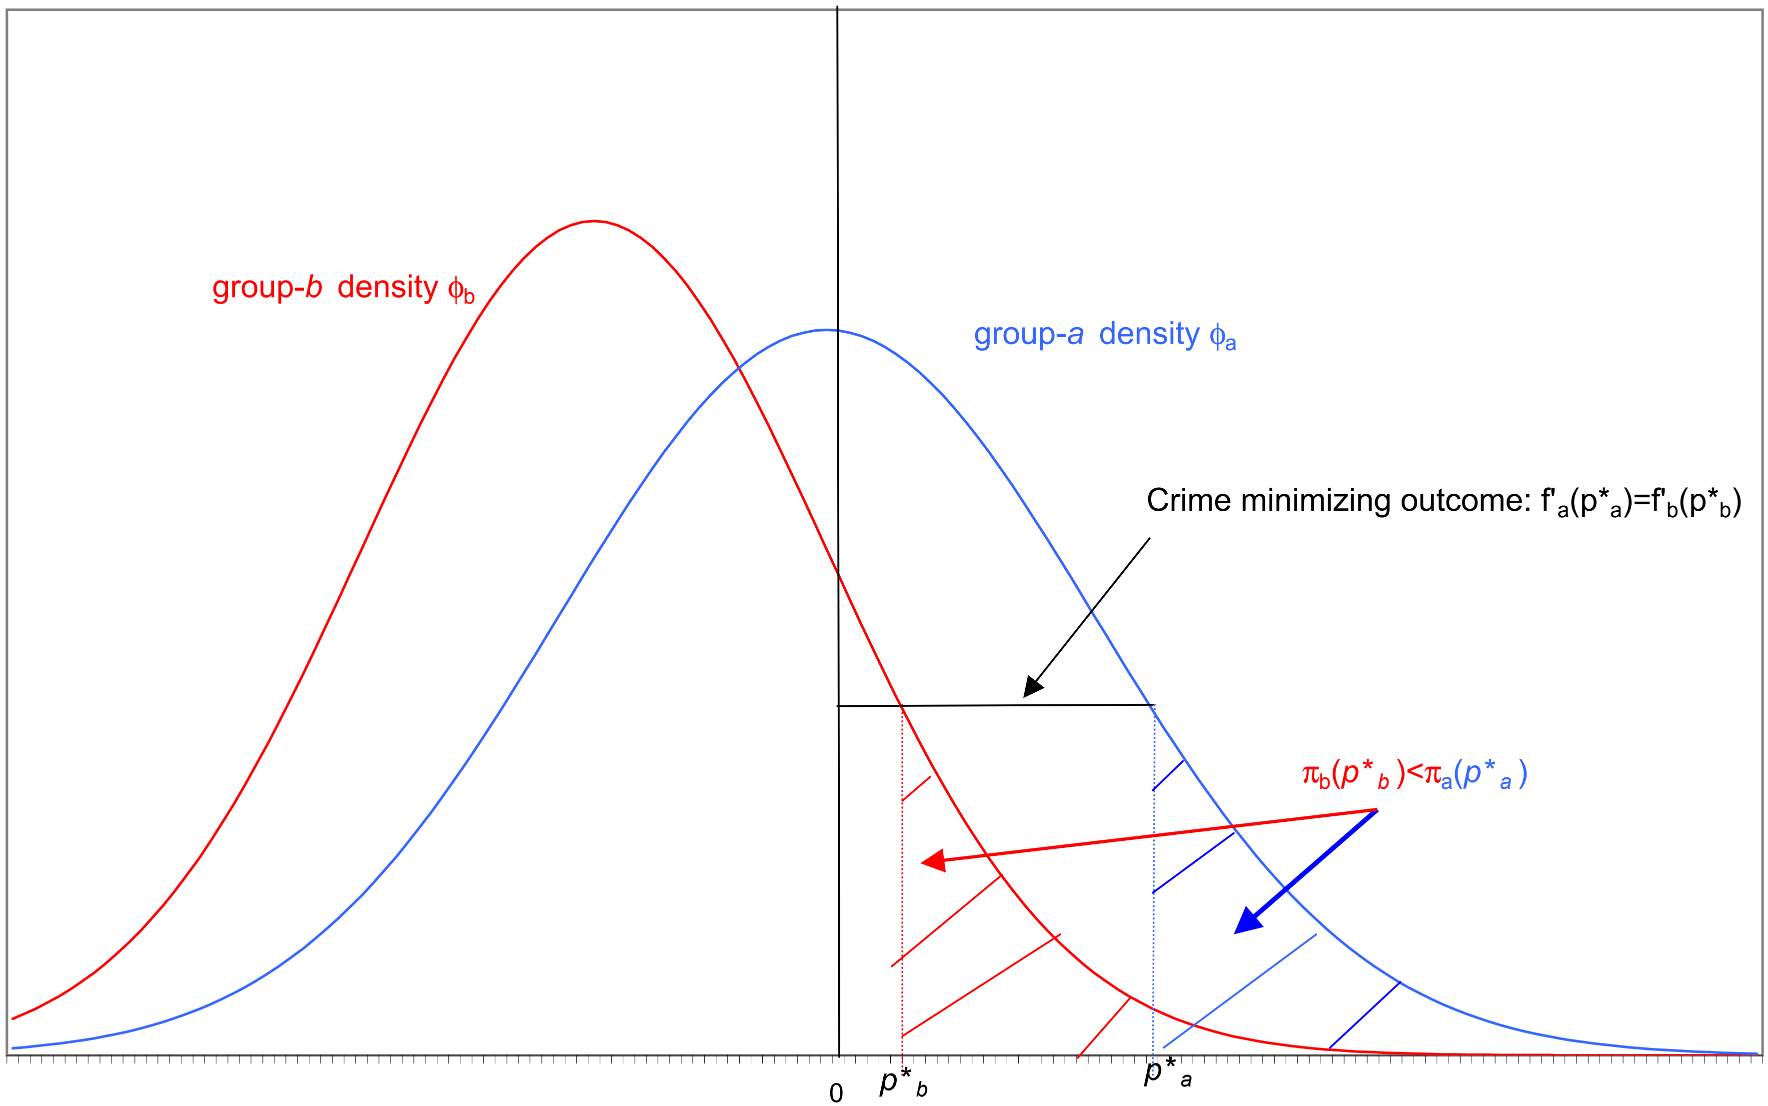

Figure 1b: Shift and Spread

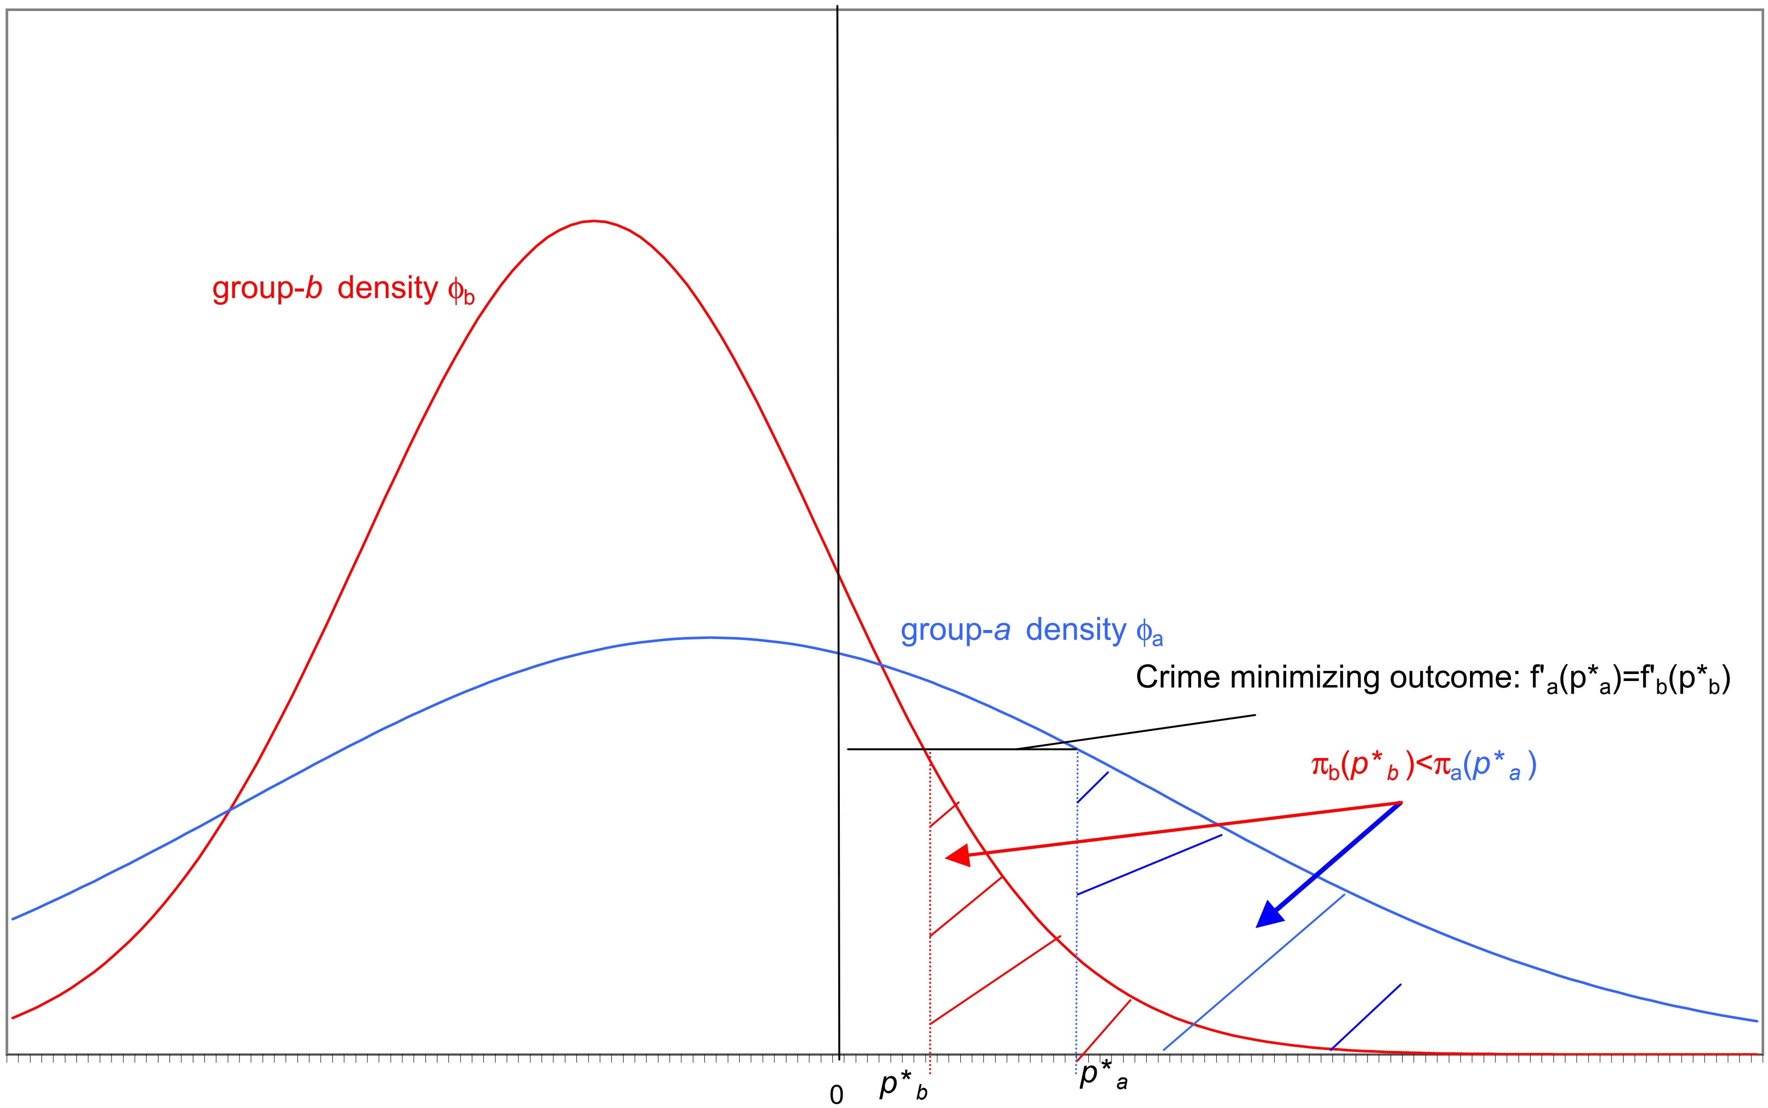

Figure 2: Shift and Tightening

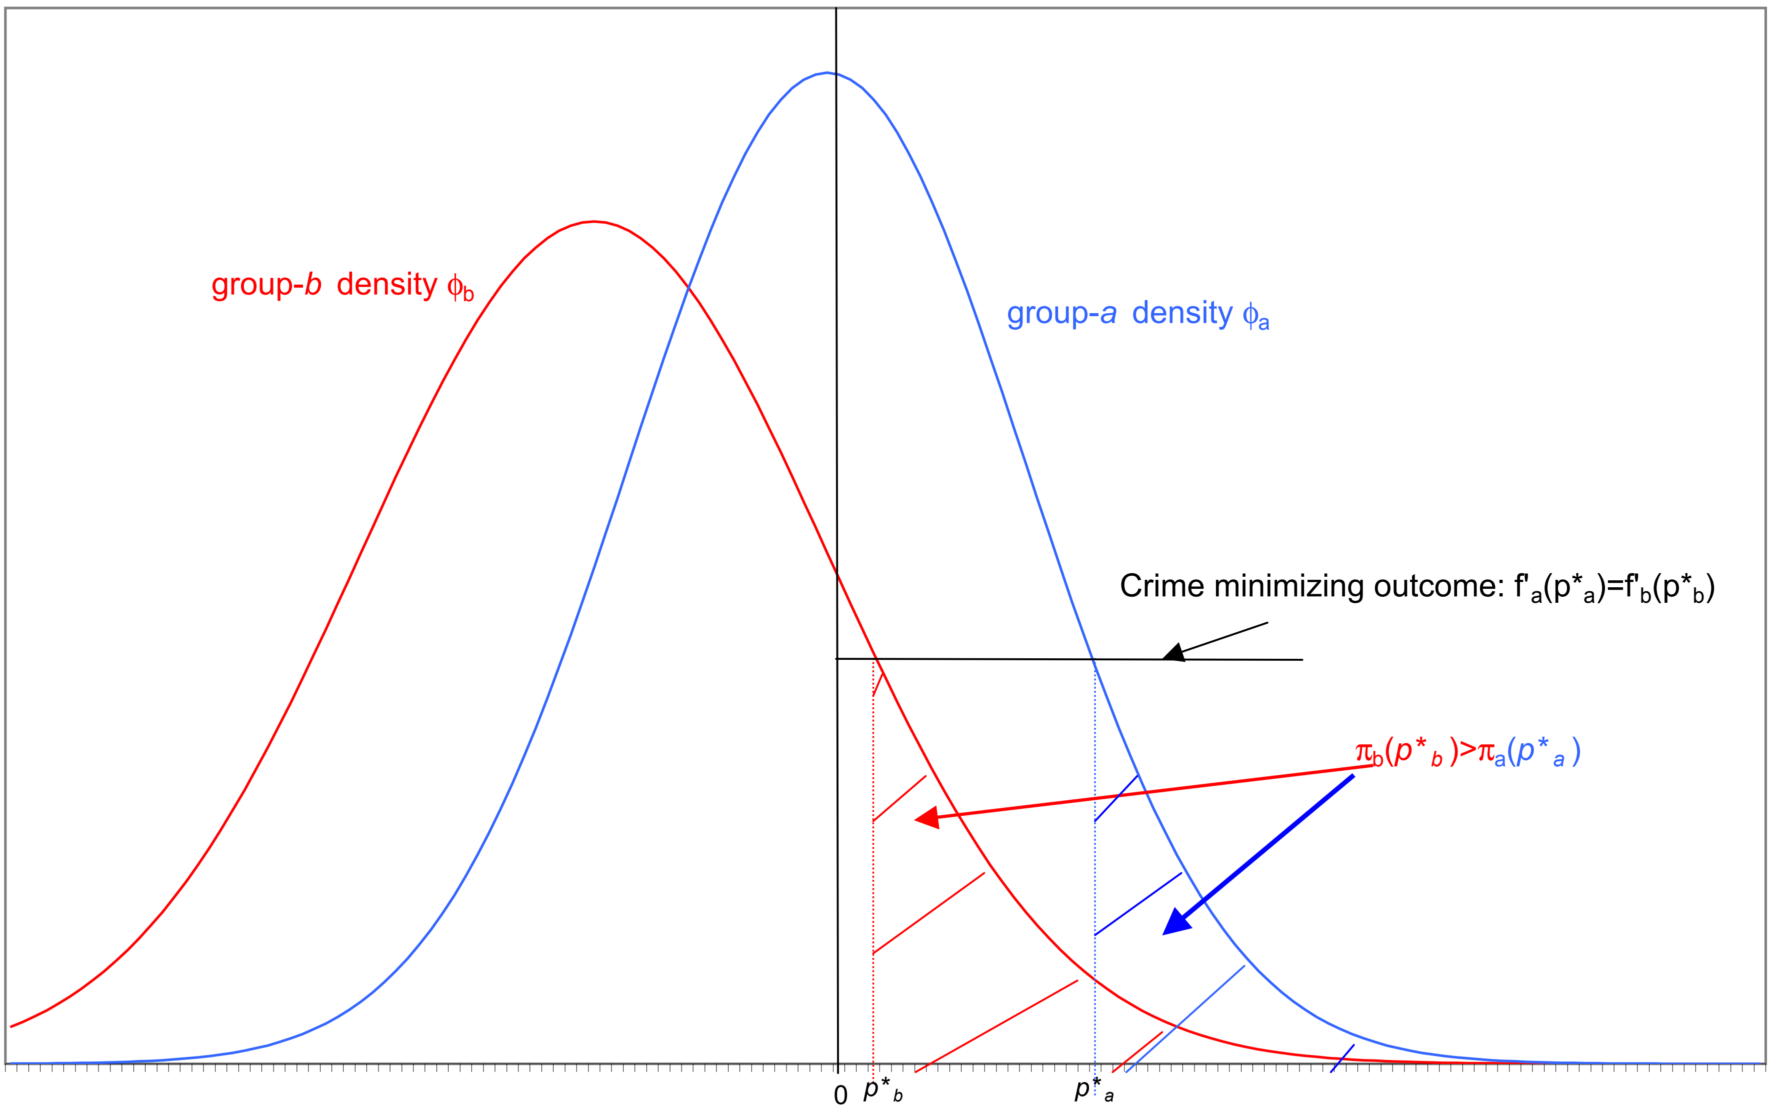

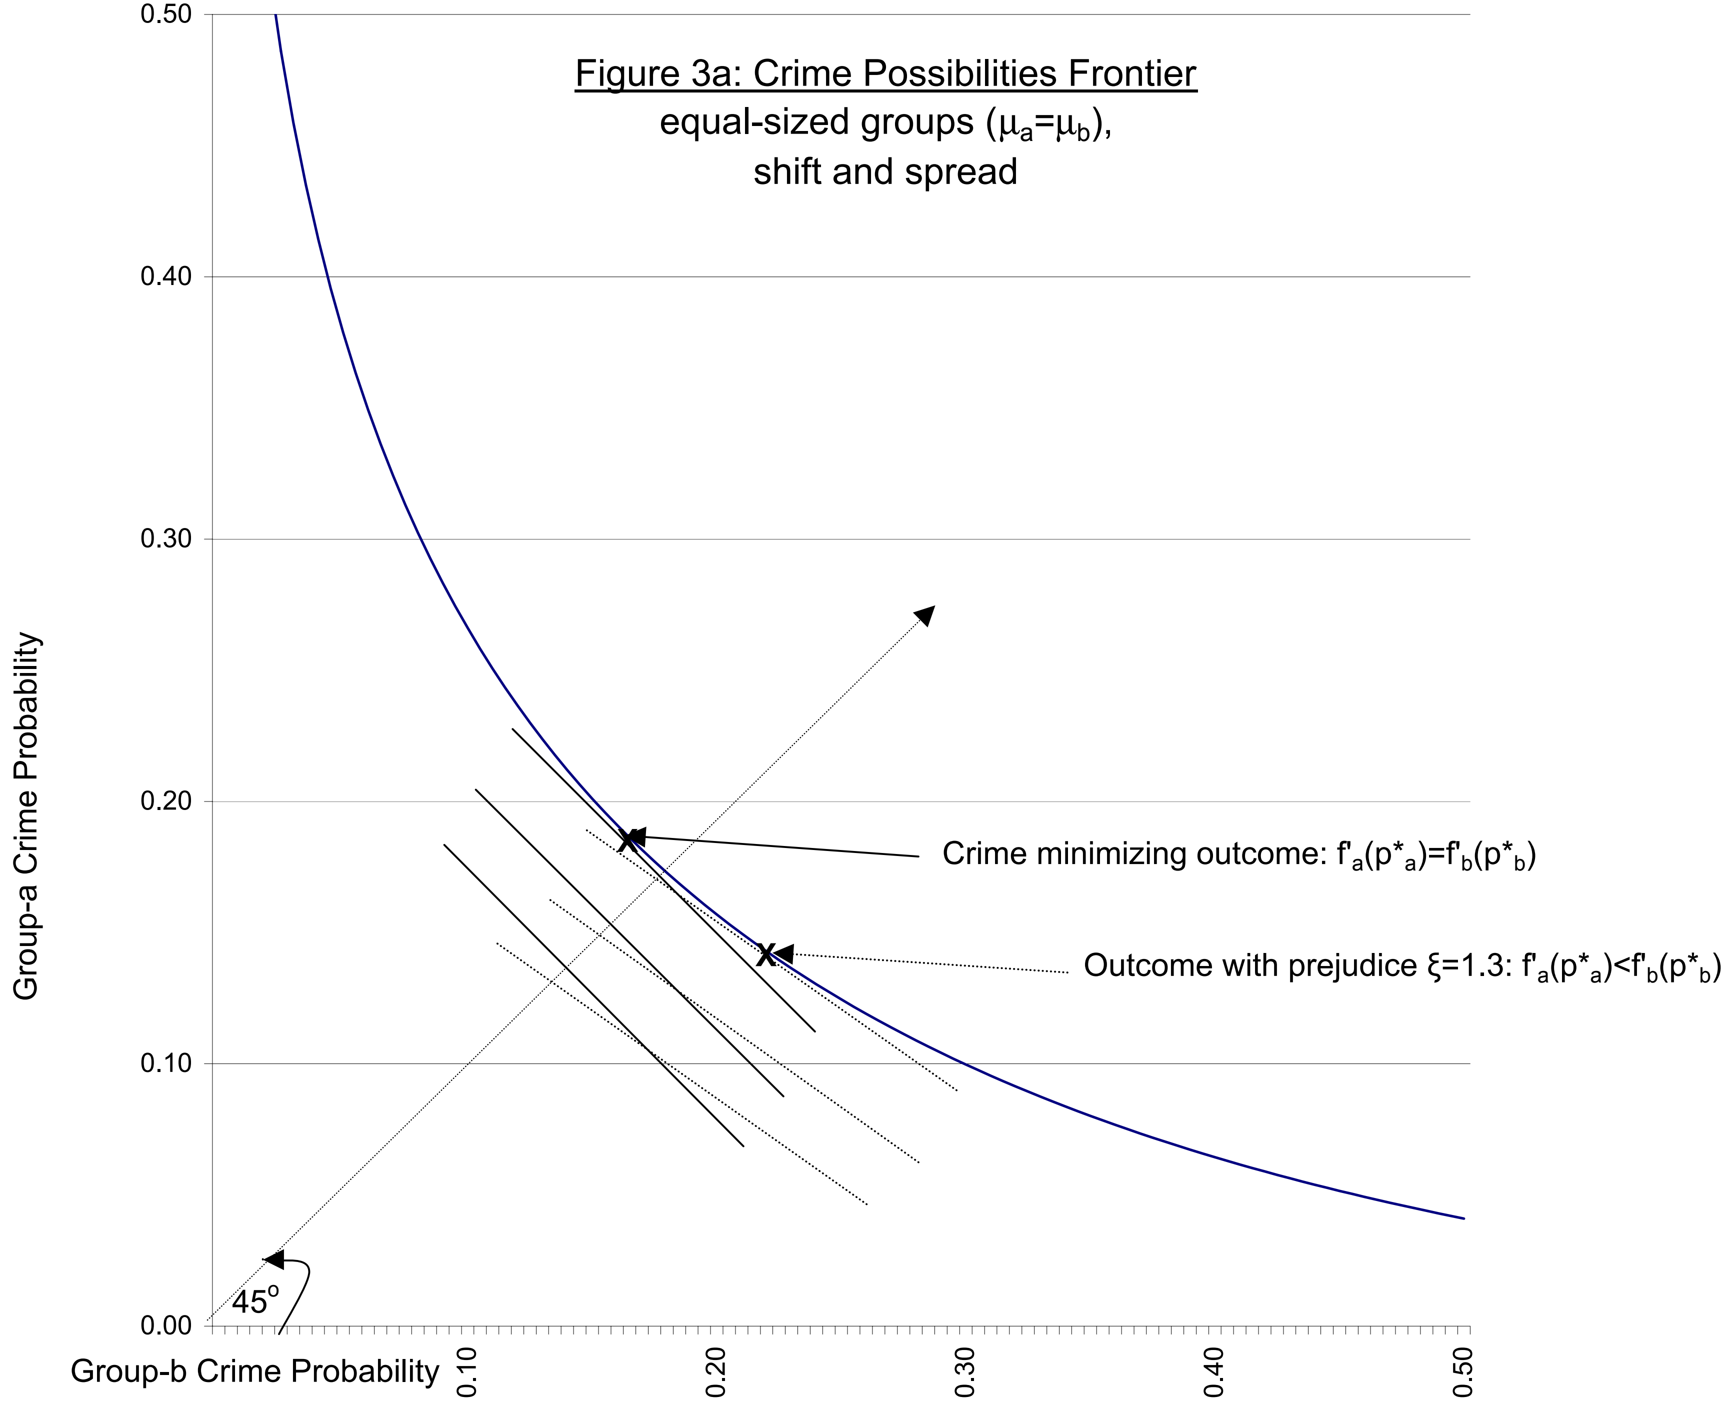

Figure 3a: Crime Possibilities Frontier
equal-sized groups ($\mu_a=\mu_b$),
shift and spread
Crime minimizing outcome: $f'_a(p^*_a)=f'_b(p^*_b)$
Outcome with prejudice ξ=1.3: $f'_a(p^*_a)<f'_b(p^*_b)$
45°
Group-a Crime Probability
Group-b Crime Probability
0.50
0.40
0.30
0.20
0.10
0.00
0.10
0.20
0.30
0.40
0.50

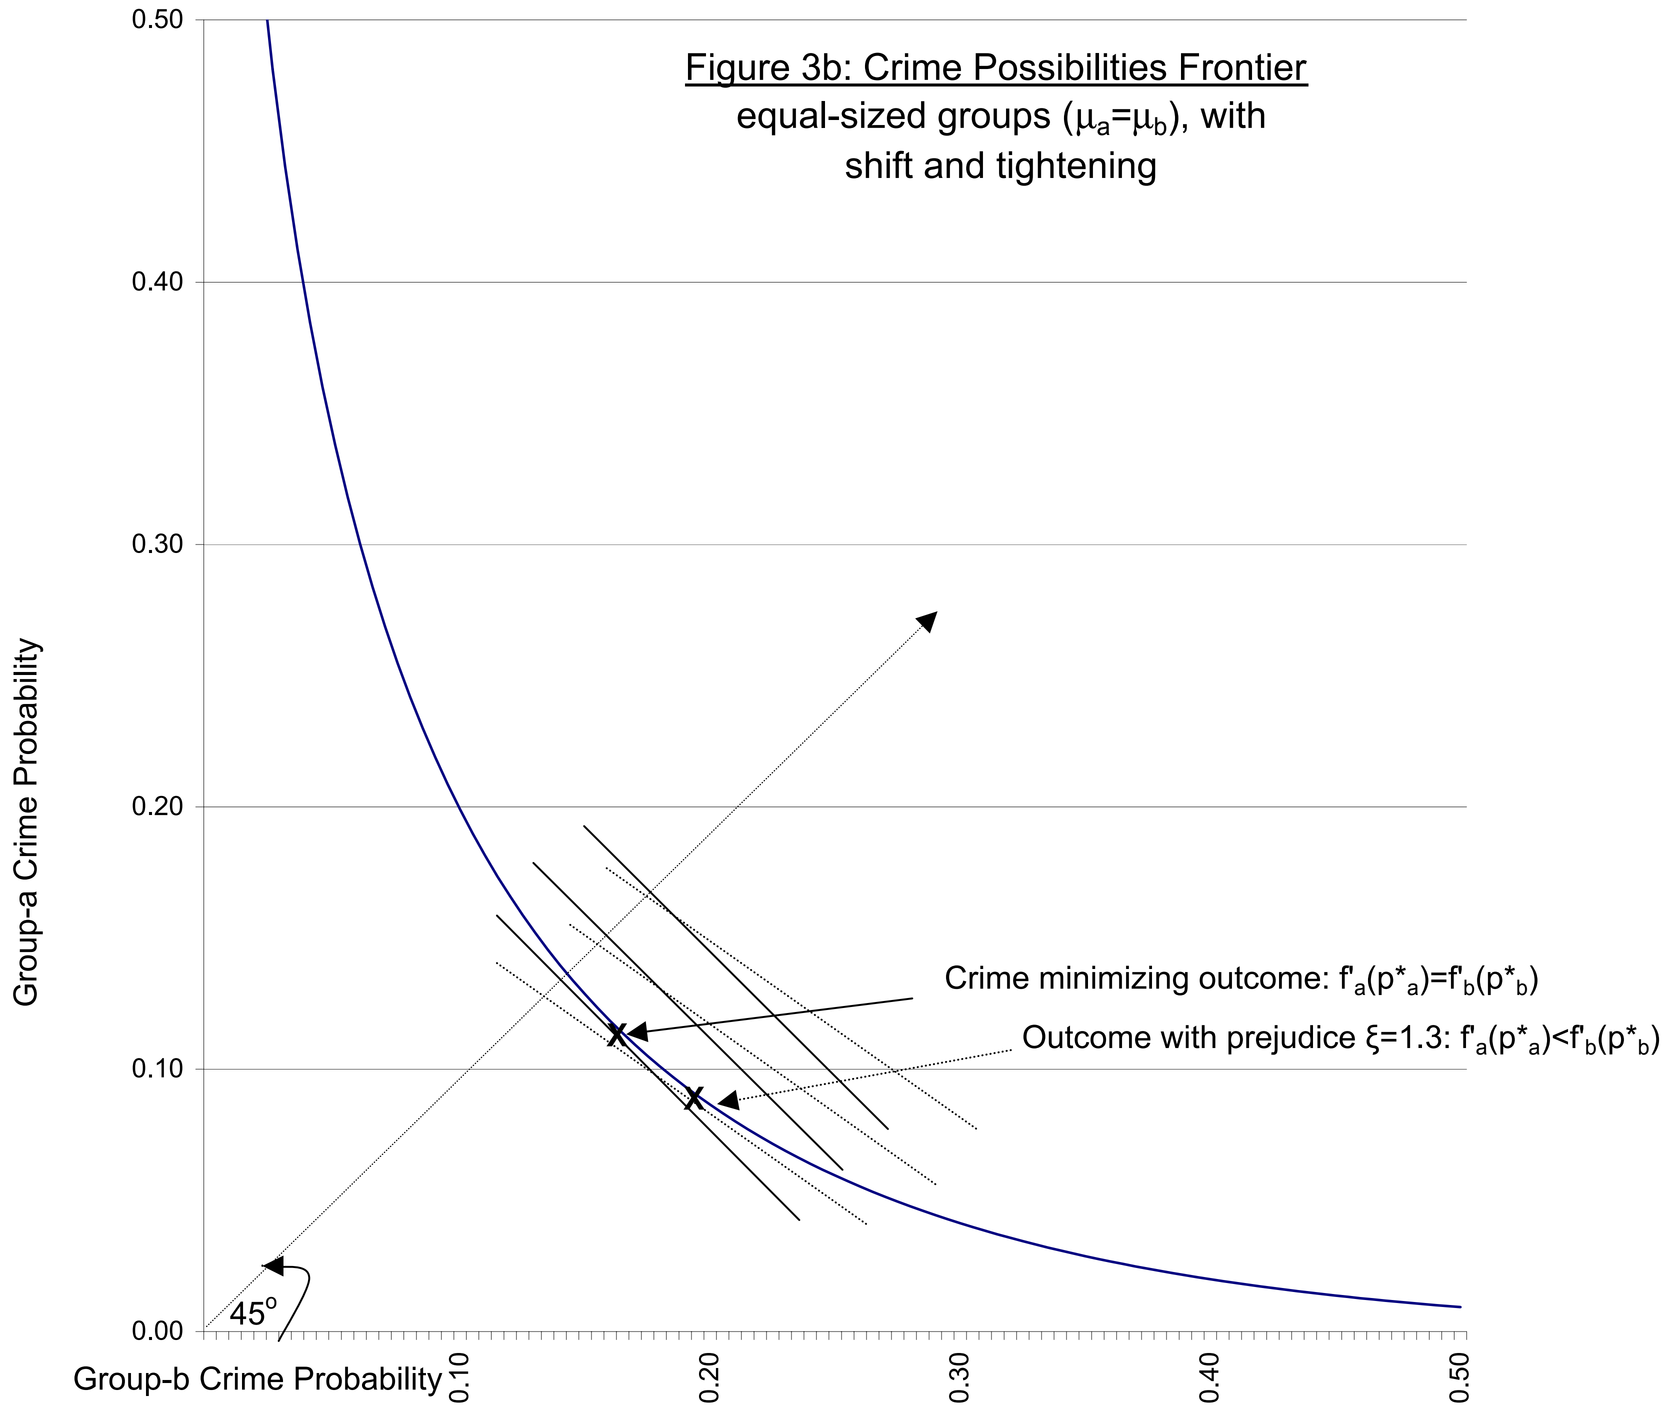
Figure 3b: Crime Possibilities Frontier
equal-sized groups ($\mu_a=\mu_b$), with
shift and tightening
Crime minimizing outcome: $f'_a(p^*_a)=f'_b(p^*_b)$
Outcome with prejudice $\xi=1.3$: $f'_a(p^*_a)<f'_b(p^*_b)$
45°
Group-a Crime Probability
Group-b Crime Probability
0.00
0.10
0.20
0.30
0.40
0.50

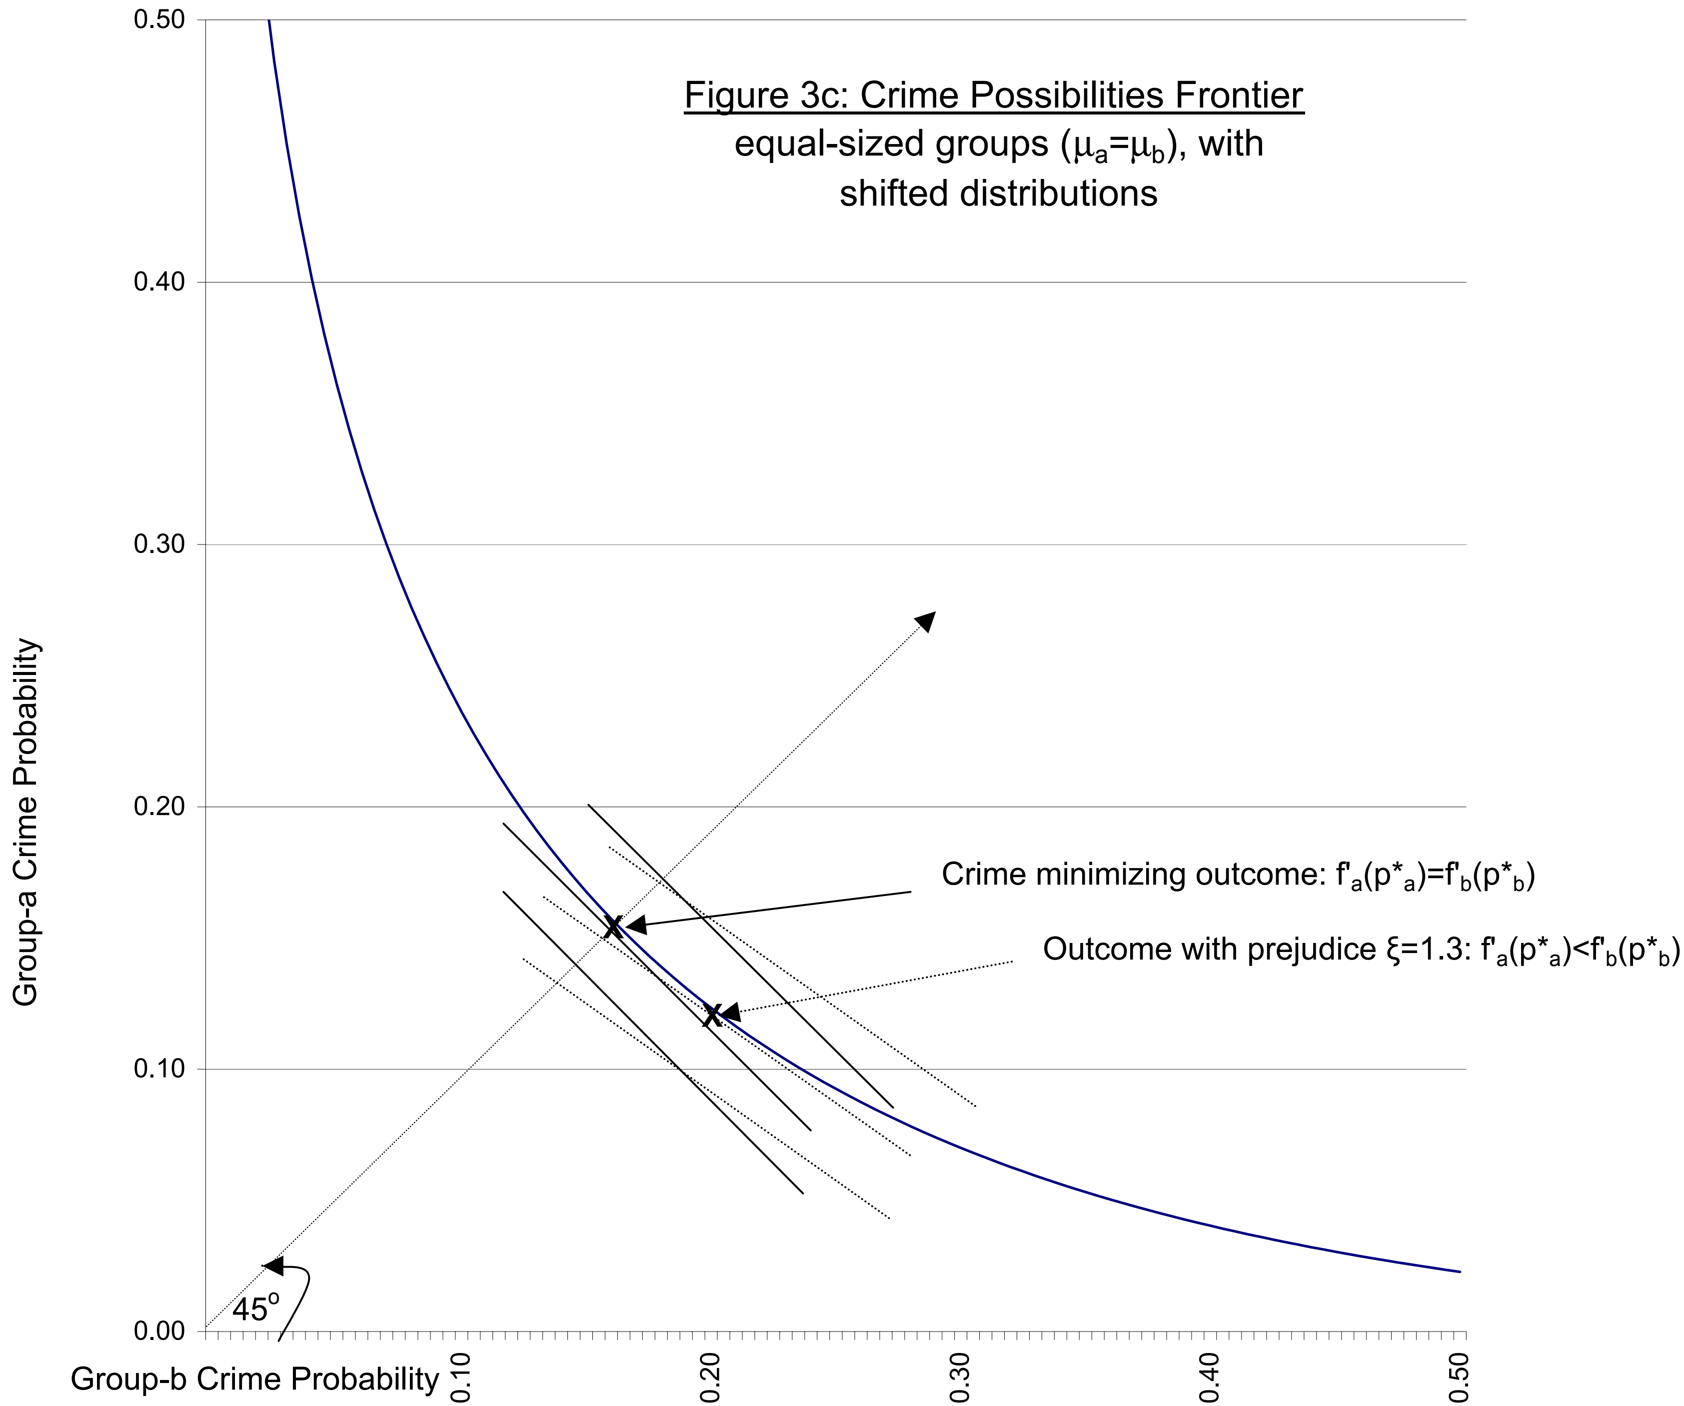
Figure 3c: Crime Possibilities Frontier
equal-sized groups (μa=μb), with
shifted distributions
Crime minimizing outcome: f'a(p*a)=f'b(p*b)
Outcome with prejudice ξ=1.3: f'a(p*a)<f'b(p*b)
45°
Group-a Crime Probability
Group-b Crime Probability
0.00
0.10
0.20
0.30
0.40
0.50