

A Behavioral Theory of Discrimination in Policing*

Ryan Hübert[†] Andrew T. Little[‡]

October 2021

Abstract

A large economic literature studies whether racial disparities in policing are explained by animus or by beliefs about group crime rates. But what if these beliefs are incorrect? We analyze a model where officers form beliefs using crime statistics, but don't properly account for the fact that they will detect more crime in more heavily policed communities. This creates a feedback loop where officers over-police groups that they (incorrectly) believe exhibit high crime rates. This inferential mistake can exacerbate discrimination even among officers with no animus and who sincerely believe disparities are driven by real differences in crime rates.

Keywords: policing, discrimination, misspecified models

JEL Classification: D83, D91, J15, K42

Declarations of interest for Ryan Hübert: none

Declarations of interest for Andrew T. Little: none

6,209 words with 2 exhibits

*Authors are listed in alphabetical order. Many thanks to Katie Angell, Tony Carey, Tom Clark, Brandon Davis, Peter Enns, Jack Glaser, Becca Goldstein, Anna Harvey, Kris Kanthak, Brad Leveck, Meg Meyer, Emily Nix, Kyle Peyton, Tessa Provins, Gwen Prowse, Marcel Roman, Will Spaniel, Allison Stashko, Jeff Staton, Emily West, Michelle Wier, Jon Woon, and audience members at the 2019 Political Economy and Public Law Conference, APSA 2019, SPSA 2020, the 2021 Conference on New Approaches to Discrimination Research, Berkeley, Chicago, Oxford, Pittsburgh, Stanford, Wash U, and USC for helpful comments.

[†]Assistant Professor, Department of Political Science, University of California, Davis.
Email: rhubert@ucdavis.edu. *Website:* <http://www.ryanhubert.com/>. Corresponding author.

[‡]Associate Professor, Department of Political Science, University of California, Berkeley.
Email: andrew.little@berkeley.edu. *Website:* <http://www.andrewtlittle.com/>.

Racial disparities in policing are pervasive. For example, 80% of people stopped under New York City’s now-defunct “stop and frisk” policy were either Black or Latino despite the fact that those two groups make up only half of the city’s population (Goel, Rao, and Shroff 2016). In Boston, Black residents comprised 63% of police stops that did not end in arrest from 2007 to 2010, even though only 24% of the population is Black (The Sentencing Project 2015). Non-white motorists are more likely to be stopped than white motorists (Epp, Maynard-Moody, and Haider-Markel 2014).

There are two standard theoretical explanations for these disparities, one driven by preferences and one driven by beliefs. In a purely preference-driven account—often called taste-based discrimination—officers intrinsically like being punitive towards some groups, or dislike being punitive towards others. The second explanation—typically called statistical discrimination—is that there are real differences in the rates of criminal behavior across groups. Knowing this, police allocate more time policing members of groups with higher crime rates, or at least in geographical areas where those groups are concentrated.

Another explanation for policing disparities sits uncomfortably between these two standard explanations. What if officers police a certain group more intensely because they believe that the group has a relatively high crime rate, but this belief is incorrect, or at least exaggerated?¹ In a proximate sense, this is discrimination driven by beliefs. But we might suspect that such inaccurate beliefs are more likely to be held by those with an intrinsic dislike of the group. If so, it makes less sense to think of the belief and preference channels as distinct and separable causes of discrimination. Instead, they may be fundamentally intertwined.

This is not just a hypothetical. Even when highly-trained researchers use administrative data, it is difficult to correct for—or even know the extent of—statistical problems (Heckman and Durlauf 2020; Knox, Lowe, and Mummolo 2020; Knox and Mummolo 2020). There is no obvious reason to think that police decision-makers will generally do better when interpreting crime statistics

¹As discussed in more detail below, several recent papers consider this possibility and provide empirical tests, though not in the context of policing (Bohren et al. 2019; Bohren, Imas, and Rosenberg 2019; Mengel and Campos Mercade 2021).

(see, e.g., Glaser 2015, for an overview). Police officials typically need to make decisions under time pressure without the benefit of the kind of statistical expertise that would enable high quality assessments about crime across communities. If departments rely on bad data or don't interpret it correctly, this can perpetuate disparities (see, for example, Harcourt 2007; Lum and Isaac 2016). Even the federal courts have weighed in to criticize flawed data analysis by police (e.g., *Floyd v. New York*, 959 F. Supp. 2d 540, S.D.N.Y. 2013).

A growing theoretical literature on agents with “misspecified models” provides a natural way to explore the implications of officers not interpreting data correctly (Esponda and Pouzo 2016; Heidhues, Koszegi, and Strack 2018; Bohren 2016). We build on a strand within this literature on how incorrect beliefs and behavior can interact (see Esponda and Pouzo 2016, for a general analysis of such games). For example, Levy and Razin (2017) study how labor market discrimination against those attending public schools and incorrect beliefs about the productivity of those attending public school can coevolve when those who attend private schools tend to be systematically pessimistic about public schools and don't adjust for this fact when learning from each other.²

We develop a model of policing where officers have a misspecified model in the sense that they do not fully account for the fact that more crimes are detected among members of groups that are policed more intensely (for previous discussions of this mechanism see Glaser 2006, 2015). When officers are deciding how to allocate resources across two communities, this misspecified model ends up inducing them to hold incorrect beliefs about the *relative* prevalence of crime among members of those two communities. We call this *non-conditioning bias*.

Our model also allows for police officers to have racial animus, and for crime rates to be different across groups. In the special case where officers form correct beliefs, these two mechanisms independently affect policing disparities, as in the standard accounts. However, once officers exhibit any non-conditioning bias, this creates a feedback channel where groups who are policed

²Outside of the study of discrimination, Heidhues, Koszegi, and Strack (2018) show that overconfidence may cause decision-makers to form incorrect beliefs about other aspects of the world (e.g., the ability of their subordinates), which in turn changes their behavior and future inferences, potentially leading to large distortions of both. In a similar vein, Levy, Razin, and Young (2020) study a model of political competition where some voters have a misspecified model of how policies map to outcomes, and learn from the equilibrium choices made by politicians.

more intensely are viewed as having higher crime rates than they really do. This feedback loop amplifies whatever policing disparities would exist in the absence of non-conditioning bias. A taste for discrimination causes inaccurate statistical discrimination. We first formalize this argument in a simple and analytically tractable model with just one officer. Next, we extend the analysis to include multiple officers and demonstrate how the discriminatory behavior of one officer can spill over and cause others to discriminate too.

One straightforward implication of our model is that faulty data analysis by police departments may unwittingly exacerbate disparities. For example, if departments use data-driven algorithms to predict where crime is likely to occur (Collins 2018), the predictions generated by these algorithms may be highly discriminatory if they are based on simple counts of prior crimes detected by police.

Our results bolster an emerging literature suggesting that it is more difficult to empirically distinguish taste-based and statistical discrimination than standard approaches would imply (e.g., Bohren, Imas, and Rosenberg 2019; Hull 2021). In our model, an officer (or a set of officers) with no animus may discriminate against one group more than can be explained by the real facts on the ground. A researcher examining policing data generated by this officer might observe that some amount of the officer’s discrimination is “unexplained.” Often, this residual discrimination is presumed to be taste-based discrimination, even though in this context it would be a form of inaccurate statistical discrimination. So, if researchers seek to empirically isolate taste-based discrimination, standard tests may need modification (or additional, potentially stringent, assumptions) to distinguish taste-based discrimination from inaccurate statistical discrimination.

1 Model of a Single Officer

We start with a model of a single police officer (pronoun “he”), who we primarily interpret as a high-level official who makes decisions for the department as a whole, such as the chief of police. Our model is intended to study discrimination in broad decisions about how to allocate police resources, rather than individual decisions of officers to initiate an interaction with a citizen or escalate to using force (e.g., Knowles, Persico, and Todd 2001; Feigenberg and Miller, n.d.;

Hull 2021). More specifically, the officer makes a choice about how to allocate resources toward policing two groups, A and B . To avoid having to make normalizations by group size, assume the two groups are equally numerous.

The officer has a unit of resources, which we primarily interpret as time, to allocate between policing the two groups. Let w_A represent the share of time spent policing group A , with $w_B = 1 - w_A$ left for group B . We assume that the officer can choose to allocate his time evenly between the two groups, but can also choose to police on group more than the other. However, the officer can't choose to allocate *all* of his time to one group or the other. Formally, the officer chooses $w_A \in [\underline{w}, \bar{w}]$, where $0 < \underline{w} \leq 1/2 \leq \bar{w} < 1$.

In the United States, it is typically illegal for governments (including police departments) to target individuals solely on the basis of their social grouping, such as their race, religion, gender, etc. Thus, one way to think about the choice in our model is that the police department decides to target resources toward different geographical locations, which due to residential segregation, have different proportions of the two groups. In Appendix A, we provide a microfoundation for the officer's choice in which the officer decides how to allocate time between neighborhoods, and not between social groups.

We assume that the allocation of policing effort, as reflected by w_A , affects the detection of crime. As a result, our model more directly captures “proactive policing,” rather than “reactive policing” where officers respond to reports of crimes in progress or which have already occurred (e.g., via 911 calls). The model is also less applicable for crimes that are universally (or near universally) reported, such as murder. More generally, what matters for our argument is that police detect more crime among groups that commit crimes at a higher rate, and where they spend more resources policing.

Formally, we let the amount of crime caught among members of group J be $c_J = p_J w_J$, where $p_J > 0$. The simplest way to interpret this is that p_J represents the average number of crimes committed by members of group J per unit of time, and w_J represents how much time is spent policing this group. This is the data that the officer uses to determine how to allocate his time.

While this formulation will prove particularly tractable, in the appendix we consider two important extensions. First, in Appendix F we allow the group crime rate to be decreasing in the amount of time spent policing that group, which could reflect policing having a deterrent effect on crime. In Appendix G, we consider a more abstract formulation of this possibility by allowing the number of crimes caught to be a general increasing function of $p_J w_J$, which complicates the interpretation of the parameters, but does not fundamentally change our argument. (Though see Feigenberg and Miller, n.d., who find that contraband detected is approximately linear in the number of traffic stops.)

Preferences We assume that the objective of the officer is to catch crimes. To capture the notion that the officer might have a taste for discrimination, we allow him to prefer catching crimes among one group or the other. We also assume that there are diminishing returns to the amount of crime caught within each group. This is a reduced-form way to capture the notion that some crimes are “more important” to detect than others, and that the officer will first dedicate time to detecting the more important crimes (within each group). In addition to these key assumptions, we place several technical assumptions on the officer utility:

Assumption 1. *The officer utility is $u(t_A c_A, t_B c_B)$, where $t_J > 0$, and the utility function $u(x_1, x_2)$ is (i) symmetric in the two arguments ($u(x_1, x_2) = u(x_2, x_1)$), (ii) continuously differentiable, (iii) strictly increasing and concave in both arguments ($u_1 > 0, u_{11} < 0, u_2 > 0, u_{22} < 0$), additively separable ($u_{12} = 0$), and (iv) homogeneous with positive degree.*

The t_A and t_B terms represent the officer’s “taste” for catching crimes among group A and B , respectively. Given part (iii) of the assumption, higher values of t_J will make the officer value catching crimes among group J more.

The remaining assumptions are for technical convenience. The symmetry assumption implies that the labeling of the groups does not affect the analysis. Additive separability allows us to set aside indirect effects where allocating effort toward one group lowers the marginal return to policing the other group. (We also show in the proof of Lemma 1 that it is sufficient that the cross-partial

not be too positive or too negative relative to the concavity in each argument.) The homogeneity assumption is primarily to provide a convenient characterization of the optimal policing allocation.

If the officer has correct beliefs about the p_J parameters, then the optimal allocation of time is a straightforward maximization of his utility function. When the officer has correct beliefs, we say he has *full information*.

Lemma 1. *Let $r_t = t_A/t_B$ and $r_p = p_A/p_B$. Given Assumption 1, if the officer knows r_p then there is a unique w_A which maximizes $u(t_A c_A, t_B c_B)$, which we write as $w_A^{br}(r_t, r_p)$.*

(i) $w_A^{br}(r_t, r_p)$ is increasing in both arguments, and where $w_A^{br}(r_t, r_p)$ is interior, it is strictly increasing in both arguments.

(ii) $w_A^{br}(1, 1) = 1/2$.

Proof See the appendix.

The r_t parameter reflects the preference for catching crimes among group A , relative to group B , which captures the possibility for taste-based discrimination. We say that if $r_t > 1$ the officer has *animus* towards group A , and $r_t < 1$ indicates animus towards group B . The r_p parameter reflects the *true* ratio of the two groups' crime rates, capturing the possibility for statistical discrimination. That is, if $r_p > 1$, the crime rate among members of group A is higher than the crime rate among members of group B , and if $r_p < 1$, the opposite is true. In the absence any asymmetry on these parameters—i.e., $r_t = 1$ and $r_p = 1$ —then part (ii) of the lemma shows that the officer splits his time equally between policing the two groups.

Importantly, all of our following results will hold for any utility function with the properties of Lemma 1. So, for example, the officer does not necessarily need to be motivated by maximizing the amount of crime caught; they could also care about prevent crime from happening in the first place (see Stashko 2020, for a an approach to disentangling these motives). What really matters is that they want to allocate more time policing groups with higher crime rates, as well as groups against which they have animus. In Appendix F, we show that these properties also hold if the crime rate is endogenous to the officer allocation and if deterrence is the officer's goal.

To reduce the cases to consider, we place one more assumption on the utility function which ensures an interior solution with full information:

Assumption 2.

$$\frac{u_2(\underline{w}, 1 - \underline{w})}{u_1(1 - \underline{w}, \underline{w})} < r_t r_p < \frac{u_2(\bar{w}, 1 - \bar{w})}{u_1(1 - \bar{w}, \bar{w})}$$

In words, this states that at the lower bound \underline{w} the marginal return to policing group A is higher than the marginal return to group B , and at the upper bound the reverse is true.

Lemma 2. *Given Assumptions 1 and 2, the full information policing allocation is interior: $w_A^{br}(r_t, r_p) \in (\underline{w}, \bar{w})$.*

As we will show, this does not preclude a corner solution when we allow the officer to have incorrect beliefs. That is, to highlight the effect of incorrect belief formation, we want to start at a benchmark where the officer chooses to spend more than the minimal amount of time policing both groups.

Main Example For illustrations, we will use the following utility function which meets assumption 1 and leads to tidy closed form solutions:

$$u(c_A, c_B) = \sqrt{t_A c_A} + \sqrt{t_B c_B} = \sqrt{t_A p_A w_A} + \sqrt{t_B p_B (1 - w_A)} \quad (1)$$

With this utility, the officer optimal allocation as a function of the ratios is:

$$w_A^{br}(r_t, r_p) = \frac{r_t r_p}{1 + r_t r_p}$$

Note that since $r_t > 0$ and $r_p > 0$, this is always strictly between 0 and 1, and hence as long as \underline{w} and \bar{w} are sufficiently close to 0 and 1, then $w_A^{br}(r_t, r_p)$ is interior.

Disparities Going forward, we now label the officer’s optimal allocation with full information as $w_A^\dagger = w_A^{\text{br}}(r_t, r_p)$. Then, whenever the officer polices one group more than the other group, there is a *policing disparity*, given by:

$$\Delta^\dagger \equiv |w_A^\dagger - 1/2|.$$

Since our model allows for both taste-based and statistical discrimination (via parameters r_t and r_p), Δ^\dagger can be decomposed into two component parts. Formally, define $w_A^{\text{stat}} = w_A^{\text{br}}(1, r_p)$ to be the “statistical policing” allocation, which reflects what an officer does if he has no animus toward either group ($r_t = 1$) but statistically discriminates based on differences in the (true) crime rates. Following Bohren et al. (2019), we refer to $w_A^{\text{stat}} - 1/2$ as the “traditional statistical discrimination,” which will contrast with the “inaccurate statistical discrimination” which arises when the officer does not know r_p . The difference between what the officer chooses and this statistical benchmark, $w_A^\dagger - w_A^{\text{stat}}$, represents taste-based discrimination. Taken together, the policing disparity when the officer has full information can be decomposed as follows:³

$$\Delta^\dagger = \underbrace{|(w_A^\dagger - w_A^{\text{stat}})|}_{\text{taste-based discrimination}} + \underbrace{|(w_A^{\text{stat}} - 1/2)|}_{\text{traditional statistical discrimination}} = |w_A^\dagger - 1/2|$$

1.1 Policing with a Misspecified Model

We now turn to our main analysis, which considers a situation in which the officer does not know the relative crime rates of the two groups (r_p), and forms this belief based on data generated by his policing choices. For now, we assume that the data the officer in our model uses is entirely driven by the crime detected as a result of his policing choices. In Section 2, we explore the consequences of officers making choices based on data generated by *other* officers’ choices as well.

³Note that taste-based and statistical discrimination may yield discrimination against different groups. In this case, the policing disparity under full information will be closer to zero than the disparities generated by either kind of discrimination on its own.

We assume that the officer may use a misspecified model of crime prevalence when forming his beliefs about relative crime rates such that he misunderstands how policing allocations affect the crime detected. In the extreme, the officer might infer that the relative number of crimes caught among the two groups is the same as the relative crime rates. As this involves the officer forming a posterior belief without properly conditioning on all relevant information, it is *non-conditioning bias*.

In the main text, we pick a particular form of this bias which leads to tidy calculations; see section Appendix D for a more general analysis. Suppose the officer forms beliefs about the crime rates as if the crime prevalence is given by:

$$\tilde{c}(w_J, p_J) = (1 - \nu)p_J w_J + \nu p_J$$

That is, his (potentially misspecified) model of crime detection is a weighted average of the *true* amount of crime detected and the crime rate ignoring his own policing efforts.

If the officer has this model of crime detection in his head, then after observing a crime rate c_J and his own policing intensity w_J , his (possibly distorted) belief \tilde{p}_J solves:

$$c_J = (1 - \nu)\tilde{p}_J w_J + \nu \tilde{p}_J = \tilde{c}(w_J, \tilde{p}_J)$$

Since $c_J = c(w_J, p_J) = w_J p_J$, we can substitute and rearrange:

$$\tilde{p}_J = \frac{w_J p_J}{(1 - \nu)w_J + \nu}$$

If $\nu = 0$, this simplifies to p_J , and is unaffected by w_J . However, for any $\nu > 0$, the belief will increase in w_J .

Combining the group ratios the belief about the ratio is:⁴

$$\tilde{r}_p(w_A, \nu) = \frac{\frac{p_A w_A}{\nu + (1-\nu)w_A}}{\frac{p_B(1-w_A)}{\nu + (1-\nu)(1-w_A)}} = r_p \left(\frac{w_A(\nu + (1-\nu)(1-w_A))}{(1-w_A)(\nu + (1-\nu)w_A)} \right). \quad (2)$$

For conciseness, we will suppress the ν argument of \tilde{r}_p in the remainder of the analysis. As ν approaches zero, the officer’s belief about crime, \tilde{r}_p , becomes more accurate (i.e., approaches r_p). As ν approaches one, \tilde{r}_p approaches the belief formed by the most extreme non-conditioning bias. More generally, as ν increases, the officer makes a more severe inferential mistake.

Regardless of the specific mechanism that generates this belief, our assumption that police officers exhibit non-conditioning bias is not outlandish. Recent studies provide causal evidence that experimental subjects neglect selection effects and thus make faulty inferences about a state of the world (e.g., Barron, Huck, and Jehiel 2019; Enke 2020). Several examples suggest the phenomenon extends to the real-world context of policing. Using a case study of drug arrests in Oakland, California, Lum and Isaac (2016) demonstrate that data used in predictive policing algorithms perpetuates policing disparities since it is generated from past policing patterns and does not appear to reflect *actual* drug use patterns. In her opinion in *Floyd v. New York*, U.S. District Judge Scheindlin writes “The City and its highest officials believe that blacks and Hispanics should be stopped at the same rate as their proportion of the local criminal suspect population” (p. 9). This is a prime example of non-conditioning bias, which is precisely what Judge Scheindlin finds troubling: “Instead, I conclude that the benchmark used by plaintiffs’ expert—a combination of local population demographics and local crime rates (*to account for police deployment*) is the most sensible” (p. 9, emphasis added). Finally, Glaser (2015) recounts a particularly clear example of non-conditioning bias when a former Los Angeles police chief told a reporter: “if officers are looking for criminal activity, they’re going to look at the kind of people who are listed on crime reports” (p. 96). Of course, the “kinds of people who are listed on crime reports” will be disproportionately from highly policed communities and not necessarily representative of those

⁴Algebraically, our bias ends up resembling a technology used in Benabou and Tirole (2006), who use it to model how individuals bias their future beliefs by limiting recall of particular kinds of information and not fully adjusting for this limited recall.

who are prone to commit crimes.

So far, we have characterized how the officer's beliefs respond to his actions and how his actions respond to his beliefs. This suggests a natural equilibrium definition:

Definition 1. *An equilibrium of the single officer model is a policing allocation w_A^* and a belief about crime rates \tilde{r}_p^* , where*

$$(i) \ w_A^* \text{ solves } w_A^* = w_A^{br}(r_t, \tilde{r}_p^*); \text{ and}$$

$$(ii) \ \tilde{r}_p^* = \tilde{r}_p(w_A^*).$$

If $\left. \frac{\partial w_A^{br}}{\partial w_A} \right|_{w_A=w_A^*} < 1$, we say the equilibrium is stable.

This is similar to what Esponda and Pouzo (2016) call a ‘‘Berk Nash Equilibrium.’’ Much of the theoretical literature on misspecified models provides general conditions under which beliefs and behavior do in fact converge to such a stable point (Esponda and Pouzo 2016; Bohren 2016; Levy, Razin, and Young 2020). To keep the focus on our application, we will analyze behavior at a stable point.

Proposition 3. *A stable equilibrium exists in the single officer model. If ν is sufficiently small, the equilibrium is unique.*

Main Example (continued) This condition on ν is not always necessary; in fact, for our main example with a utility function given by (1), there is a unique equilibrium in which the officer chooses a policing allocation

$$w_A^* = \begin{cases} \underline{w} & \text{if } \hat{w}_A < \underline{w} \\ \hat{w}_A & \text{if } \hat{w}_A \in [\underline{w}, \bar{w}] \\ \bar{w} & \text{if } \hat{w}_A > \bar{w} \end{cases}$$

where

$$\hat{w}_A = w_A^\dagger + \frac{\nu(r_t r_p - 1)}{(1 - \nu)(1 + r_t r_p)}$$

This is because there is a unique solution to $w_A = \frac{r_t \tilde{r}_p(w_A)}{1+r_t \tilde{r}_p(w_A)}$ given by \hat{w}_A . If \hat{w}_A lies in $[\underline{w}, \bar{w}]$, then it corresponds to an equilibrium allocation. Whenever \hat{w}_A does not lie in $[\underline{w}, \bar{w}]$, then there is an equilibrium at a corner solution.⁵

To allow for clean statements about how inaccurate beliefs affect equilibrium behavior, our remaining technical results (and illustrations) focus on the case where there is a unique equilibrium.

Illustrations In each panel of Figure 1, we illustrate the fixed point analysis for the officer’s decision problem, using different values of r_t and r_p and assuming a utility function given by (1). The black curves trace out $w_A^{\text{br}}(r_t, \tilde{r}_p(w_A))$ as a function of w_A . The grey 45-degree line represents points where the best response allocation equals the actual allocation. Starting at any point w_A , if the $w_A^{\text{br}}(r_t, \tilde{r}_p(w_A))$ curve lies above the 45 degree line, then an officer who initially policies at allocation w_A will generate a belief about the relative crime rates that makes him want to police group A more. Conversely, if the curve lies below the 45 degree line, an officer starting at w_A would want to police group A less. An equilibrium allocation lies at an intersection of the black curve and the 45 degree line.

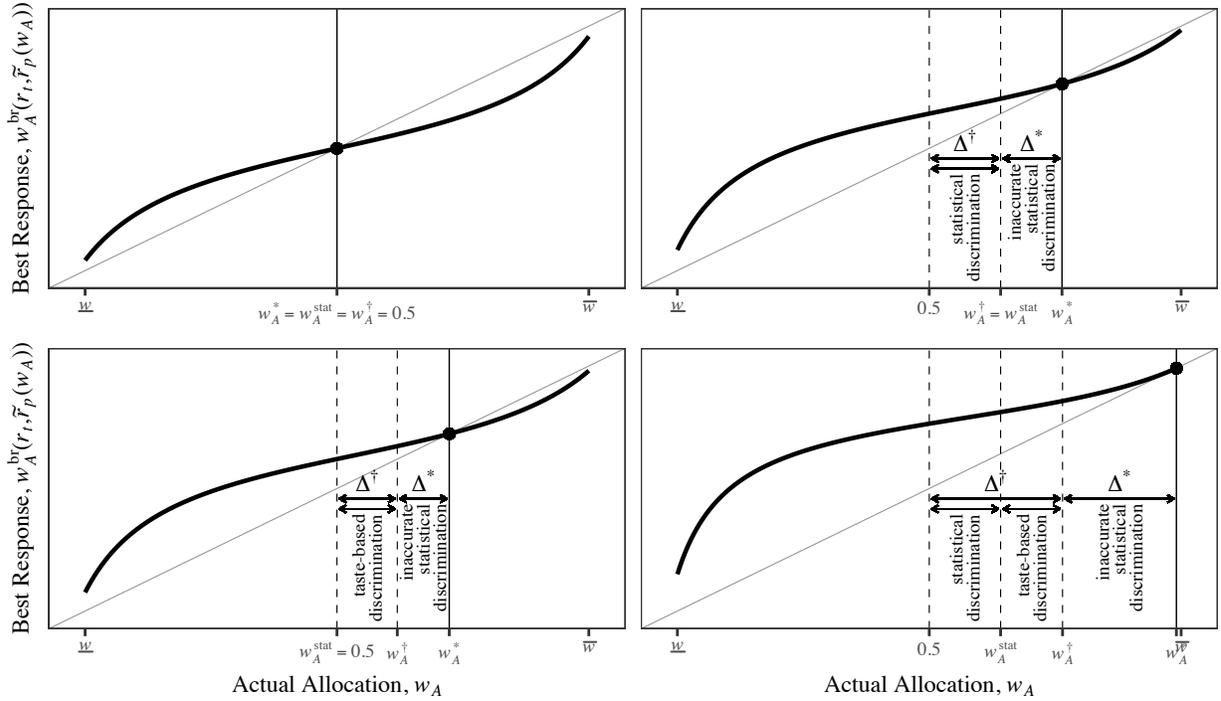
The difference between the panels in Figure 1 is that in the top panels the officer has no animus towards either group ($r_t = 1$), while in the bottom panels he has animus towards group A . In the left panels there is no difference in actual crime rates ($r_p = 1$), while in the right panels the true crime rate is higher in group A ($r_p > 1$).

The top left panel depicts a scenario with equal crime rates and no officer animus, $r_p = r_t = 1$. In this situation, despite making inferential mistakes, the officer’s policing allocation is equal, $w_A^* = 1/2$. If the officer were to police group A more or less, there would be “self-correction” in the sense described above: he would move back towards the equilibrium with equal policing.

However, equal policing is fragile to changes in the exogenous parameters r_t and r_p . The bottom left panel demonstrates a situation with equal crime rates, but where the officer has animus toward group A . As the figure depicts, without making an inferential mistake, the officer’s animus toward group A causes him to engage in taste-based discrimination against group A so that $w_A^\dagger >$

⁵We formally state this equilibrium in Proposition 6 in the appendix.

Figure 1: In each panel, we plot the officer’s best response $w_A^{\text{br}}(r_t, \tilde{r}_p(w_A))$ as a function of his actual policing allocation w_A . an equilibrium of the model occurs where $w_A^{\text{br}}(r_t, \tilde{r}_p(w_A))$ intersects the diagonal line—i.e., at a fixed point, denoted by a large dot. Each panel depicts equilibria for different parameter values. We also depict the disparities caused by statistical, taste-based and inaccurate statistical discrimination in each equilibrium. For the left panels, crime rates are equal ($r_p = 1$) and for the right panels, group A ’s crime rate is higher ($r_p > 1$). For the top panels, the officer has no animus ($r_t = 1$) and for the bottom panels, the officer has animus against A ($r_t > 1$).



1/2 (and thus $\Delta^\dagger > 0$). However, his non-conditioning bias causes him to police group A even more than he would due to his animus alone, $w_A^* > w_A^\dagger$.

Formally, if the officer chooses a policing allocation w_A^* in an equilibrium, then we define the policing disparity relative to the full information benchmark as:

$$\Delta^* \equiv |w_A^* - w_A^\dagger|.$$

This is the “excess disparity” caused by the fact that the officer makes an inferential mistake when forming his belief about the two crime rates. Following (Bohren et al. 2019), we refer to it as *inaccurate statistical discrimination*. We will show below that inaccurate statistical discrimination

always goes in the same direction as the disparity caused by the standard explanations (and represented by Δ^\dagger). We can therefore denote total discrimination as $\Delta = \Delta^\dagger + \Delta^*$. Returning to the bottom left panel of Figure 1, in this equilibrium about half of the officer's discrimination is driven by taste and about half is driven by non-conditioning bias.

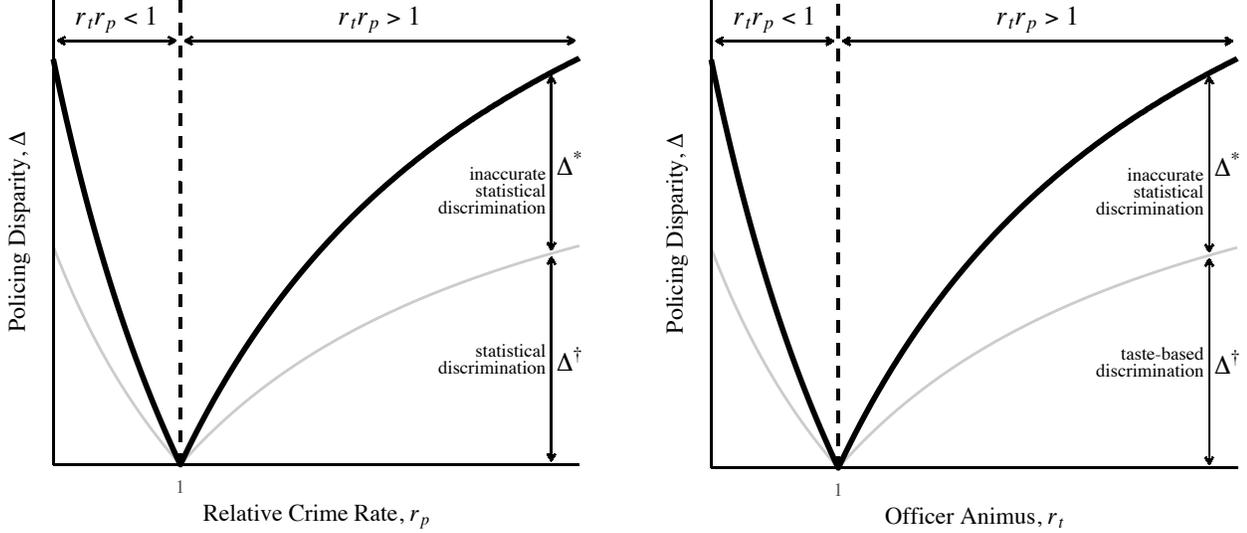
Inaccurate statistical discrimination can also occur in the absence of officer animus. The top right panel indicates a case where $r_t = 1$ but $r_p > 1$. So, some excess policing of group A is explained by different crime rates (again $w_A^\dagger > 1/2$, and $\Delta^\dagger > 0$), but the officer believes these differences are bigger than they really are. As with the illustration of taste-based discrimination, this roughly doubles the policing disparity relative to the full information benchmark. In a sense, this is all statistical discrimination, but roughly half of it is driven by false beliefs.

Finally, the bottom right panel shows a case where group A has a higher crime rate and the officer has animus towards this group. In this case, no matter what feasible allocation he chooses, he would always like to police group A even more. This leads to a corner solution even though his policing allocation would be interior if he had full information.

The officer's non-conditioning bias creates a link between taste-based and statistical discrimination. For an officer with any strictly positive level of this bias, taste-based and statistical discrimination are no longer two mutually exclusive channels through which policing disparities emerge. When conceptualized in this way, our model shows that taste-based discrimination can *cause* (inaccurate) statistical discrimination. And since an officer's animus can cause distorted beliefs about crime rates, our model maps into an intuition in the academic literature (and in popular discourse) that the empirical phenomenon of prejudice will typically involve both racial animus and incorrect beliefs.

To be more concrete about how this works in our model, consider the following. First, the officer's animus causes him to allocate more policing effort toward one group. Then, since he spends more time policing that group, he sees more crimes among members of that group. Finally, his non-conditioning bias causes him to infer that the increased number of crimes he observes is an indication that the crime rate among members of that group is higher than it actually is. As a result

Figure 2: In each panel, we plot the policing disparity that emerges in an interior equilibrium of the model, as a function of the true relative crime rate (left panel) and the officer’s animus toward group A (right panel). As long as $r_t r_p \neq 1$, the officer always engages in either statistical or taste-based discrimination, *as well as* inaccurate statistical discrimination.



(and notwithstanding his animus), his non-conditioning bias causes him to *sincerely believe* that some (or even most) his overpolicing of one group is justified by the prevalence of crime among members of that group. As Gelman, Fagan, and Kiss (2007) point out: “Police often point to the high rates of seizures of contraband, weapons, and fugitives in such stops, and also to a reduction of crime, to justify such aggressive policing” (p. 814).

The next proposition states exactly when policing with the misspecified model we study ends up amplifying policing disparities caused by taste-based and/or statistical discrimination.

Proposition 4. For any $\nu \in (0, 1)$:

- (i) If $r_t r_p = 1$, then there is an equilibrium with no policing disparity (since $w_A^* = w_A^\dagger = 1/2$), and the officer has correct beliefs about crime, $\tilde{r}_p^* = r_p$.
- (ii) If $r_t r_p \neq 1$, and the equilibrium is unique, then policing with a misspecified model amplifies existing disparities: $w_A^* > w_A^\dagger > 1/2$ if $r_t r_p > 1$ and $w_A^* < w_A^\dagger < 1/2$ if $r_t r_p < 1$ (alternatively, $\Delta^* > 0$), and the officer has incorrect beliefs, $\tilde{r}_p^* \neq r_p$.

If the officer’s policing allocation is not at a corner (\underline{w} or \bar{w}), then the disparity caused by

inaccurate statistical discrimination is strictly positive as $r_t r_p$ moves away from 1. Figure 2 illustrates. In the left panel, we plot policing disparities as a function of the (true) relative crime rate, r_p . In the right panel, we plot policing disparities as a function of the officer's animus, r_t . In each panel, the grey line depicts the policing disparity caused by statistical and taste-based discrimination and the black line depicts the entire policing disparity. Note that in either panel, as long as $r_t r_p \neq 1$, then inaccurate statistical discrimination causes the policing disparity to be higher than it otherwise would have been with only taste-based and statistical discrimination.

In this section, we have analyzed a model of policing by a representative officer, generating our core insight that, once officers have even mild forms of non-conditioning bias, taste-based discrimination can cause inaccurate statistical discrimination.

However, one potential limitation is that the data collected from the officer's own policing decisions (i.e. w_A^*) is the only thing causing him to form distorted beliefs. In reality, police departments are comprised of a multiple police officers with diverse preferences, and all of their individual policing choices end up contributing to the department's overall assessment of crime across communities. In the next section, we extend the model to look at how the presence of multiple, heterogeneous police officers affects our findings.

2 Model with Multiple, Heterogeneous Officers

To study how the dynamics of the model are different with multiple decision-makers, we analyze the simplest such environment: with two officers, indexed by $i \in \{1, 2\}$. Both officers choose how much time to allocate to group A , $w_{A,i} \in [\underline{w}, \bar{w}]$, with the remainder allocated to group B : $w_{B,i} = 1 - w_{A,i}$. We now let $w_J = w_{J,1} + w_{J,2}$ represent the *total* policing of group J , where $w_J \in [2\underline{w}, 2\bar{w}]$. Let $c_{J,i} = p_J w_{J,i}$ be the number of crimes caught among group J by officer i , and $c_J = p_J w_J$ the total crime caught among members of group J .

To simplify, we assume each officer cares only about the number of crimes that he catches, and

use the utility function of our main example in the previous section:

$$u_i(c_{A,i}, c_{B,i}) = \sqrt{t_{A,i}c_{A,i}} + \sqrt{t_{B,i}c_{B,i}} = \sqrt{t_{A,i}p_A w_{A,i}} + \sqrt{t_{B,i}p_B(1 - w_{A,i})}$$

This utility function allows us to isolate the affect of distorted beliefs on policing since it means that there is no *direct* effect of officer j 's behavior on the utility of officer i . There will only be an *indirect* effect of the other officer's behavior via officer i 's belief. If instead each officer's utility were to be defined over the total crime caught, then the policing behavior of the other officer has a direct effect on his own best response, and we would not be able to cleanly isolate how much distorted beliefs affect policing decisions.

We also define the officers' beliefs in a similar way to the single officer model, but accounting for the fact that there are now two officers making policing allocations:

$$\tilde{r}_{p,i}(w_A) = \frac{\frac{c_A}{\nu_i + (1-\nu_i)w_{A,1} + \nu_i + (1-\nu_i)w_{A,2}}}{\frac{c_B}{\nu_i + (1-\nu_i)w_{B,1} + \nu_i + (1-\nu_i)w_{B,2}}} = \frac{\frac{c_A}{2\nu_i + (1-\nu_i)w_A}}{\frac{c_B}{2\nu_i + (1-\nu_i)(2-w_A)}} \quad (3)$$

Note that each officer's belief in the multiple officer model is indexed by i since each officer can, in principle, differ with respect to the severity of their non-conditioning bias (i.e., have different values of ν_i).

In Appendix E, we consider an alternative version of this bias where the officers adjust differently for their own behavior and the other officer's behavior. The key property of the symmetric version we study here, as well as the version we study in the appendix, is that officer i 's belief about the relative prevalence of crime among members of group A is increasing in how much the *other* officer polices group A .

Formally, we define a solution of the multiple officer model as follows:

Definition 2. *An equilibrium of the model with two officers is a pair of allocation choices $(w_{A,1}^*, w_{A,2}^*)$ and vector of beliefs $(\tilde{r}_{p,1}^*, \tilde{r}_{p,2}^*)$ such that for all $i \in \{1, 2\}$:*

- (i) $w_{A,i}^* = w_A^{br}(r_{t,i}, \tilde{r}_{p,i}^*)$, and
- (ii) $\tilde{r}_{p,i}^* = \tilde{r}_{p,i}(w_A^*)$ is given by equation (3) evaluated at w_A^* .

If a condition analogous to the single officer model is met (see Appendix B), we say the equilibrium is stable.

With multiple officers, it is difficult to obtain closed-form solutions. However, it is straightforward to show that an equilibrium exists, and in any equilibrium that meets a stability condition analogous to the single-officer model (see Appendix B), an officer discriminates more when the *other* officer has more animus or higher non-conditioning bias. That is, discrimination spills over across officers.

Proposition 5. *In the model with two officers, an equilibrium exists. At any stable interior equilibrium allocation:*

- (i) *Each officer's allocation to group A ($w_{A,i}^*$) is strictly increasing in the animus of either officer, $r_{t,1}$ or $r_{t,2}$, and*
- (ii) *If the officers collectively spend more than half of their time policing group J ($w_j^* > 1$), then each officer's allocation to group J is strictly increasing in the non-conditioning bias of either officer, ν_1 or ν_2 .*

The two parts of this result demonstrate that taste-based discrimination and inferential mistakes are contagious across officers. These findings suggest that efforts to reduce policing disparities by reducing officer animus (via training), or diversifying police forces to reduce the number of officers with animus, may be of limited effectiveness as long as some officers still have animus toward one or more groups. Given their non-conditioning bias, a bad apple (or even a well-intentioned, but naive apple) can both spoil the bunch.

3 Conclusion

In this paper, we study how police officials allocate resources across two groups of citizens if they make an empirically common inferential mistake when evaluating crime rates. Departing from the standard assumption that decision-makers be fully Bayesian, we instead assume that

police officials form beliefs about the relative prevalence of crime among members of two groups without fully accounting for the intensity with which they police each of those two groups.

When combined with the possibility of either group-based animus or statistical differences across groups (or both), we show that officers with this kind of non-conditioning bias generically overpolice one of the two groups. They do so because they form exaggerated beliefs about the relative crime rate among members of that group. This *amplifies* any existing disparities caused by taste-based and/or statistical discrimination.

Our analysis suggests that when officers have non-conditioning bias, it no longer makes sense to treat taste-based and statistical discrimination as separate and independent channels through which discrimination occurs. Nor does it make sense to presume that the racial animus of an individual officer only affects that officer's behavior. Indeed, racial animus and discrimination based on incorrect beliefs are both intertwined and spill over across officers.

References

- Barron, Kai, Steffen Huck, and Philippe Jehiel. 2019. "Everyday Econometricians: Selection Neglect and Overoptimism When Learning from Others." Manuscript.
- Benabou, Roland, and Jean Tirole. 2006. "Belief in a just world and redistributive politics." *The Quarterly journal of economics* 121 (2): 699–746.
- Bohren, J. Aislinn. 2016. "Informational herding with model misspecification." *Journal of Economic Theory* 163:222–247.
- Bohren, J. Aislinn, Kareem Haggag, Alex Imas, and Devin G. Pope. 2019. "Inaccurate Statistical Discrimination," NBER Working Paper Series, no. 25935. doi:10.3386/w25935. <http://www.nber.org/papers/w25935>.
- Bohren, J. Aislinn, Alex Imas, and Michael Rosenberg. 2019. "The dynamics of discrimination: Theory and evidence." *American economic review* 109 (10): 3395–3436.
- Collins, Dave. 2018. "'Predictive Policing': Big-City Departments Face Lawsuits." *AP*. <https://apnews.com/b11e4bca11e548d3af7a63f24e348c6f>.
- Enke, Benjamin. 2020. "What You See Is All There Is." *Quarterly Journal of Economics* 135 (3): 1363–1398.
- Epp, Charles R., Steven Maynard-Moody, and Donald P. Haider-Markel. 2014. *Pulled Over: How Police Stops Define Race and Citizenship*. Chicago, IL: University of Chicago Press.
- Esponda, Ignacio, and Demian Pouzo. 2016. "Berk–Nash equilibrium: A framework for modeling agents with misspecified models." *Econometrica* 84 (3): 1093–1130.
- Feigenberg, Benjamin, and Conrad Miller. n.d. "Would Eliminating Racial Disparities in Motor Vehicle Searches Have Efficiency Costs?" Forthcoming, *Quarterly Journal of Economics*.

- Gelman, Andrew, Jeffrey Fagan, and Alex Kiss. 2007. "An Analysis of the New York City Police Department's "Stop-and-Frisk" Policy in the Context of Claims of Racial Bias." *Journal of the American Statistical Association* 102 (479): 813–823.
- Glaser, Jack. 2006. "The efficacy and effect of racial profiling: A mathematical simulation approach." *Journal of Policy Analysis and Management: The Journal of the Association for Public Policy Analysis and Management* 25 (2): 395–416.
- . 2015. *Suspect Race: Causes and Consequences of Racial Profiling*. New York: Oxford University Press.
- Goel, Sharad, Justin M. Rao, and Ravi Shroff. 2016. "Precinct or Prejudice? Understanding Racial Disparities in New York City's Stop-and-Frisk Policy." *The Annals of Applied Statistics* 10 (1): 365–394.
- Harcourt, Bernard E. 2007. *Against Prediction: Profiling, Policing, and Punishing in an Actuarial Age*. Chicago, IL: University of Chicago Press.
- Heckman, James J., and Steven N. Durlauf. 2020. "Comment on "An Empirical Analysis of Racial Differences in Police Use of Force" by Roland G. Fryer Jr." *Journal of Political Economy* 0 (ja): null. doi:10.1086/710976. eprint: <https://doi.org/10.1086/710976>. <https://doi.org/10.1086/710976>.
- Heidhues, Paul, Botond Koszegi, and Philipp Strack. 2018. "Unrealistic expectations and misguided learning." *Econometrica* 86 (4): 1159–1214.
- Hull, Peter. 2021. "What Marginal Outcome Tests Can Tell Us About Racially Biased Decision-Making." Manuscript.
- 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.

- Knox, Dean, Will Lowe, and Jonathan Mummolo. 2020. “Administrative Records Mask Racially Biased Policing.” *American Political Science Review* 114 (3): 619–637.
- Knox, Dean, and Jonathan Mummolo. 2020. “Making Inferences about Racial Disparities in Police Violence.” *PNAS* 117 (3): 1261–1262.
- Levy, Gilat, and Ronny Razin. 2017. “The coevolution of segregation, polarized beliefs, and discrimination: The case of private versus state education.” *American Economic Journal: Microeconomics* 9 (4): 141–70.
- Levy, Gilat, Ronny Razin, and Alwyn Young. 2020. “Misspecified Politics and the Recurrence of Populism.” Manuscript.
- Lum, Kristian, and William Isaac. 2016. “To Predict and Serve?” *Significance*: 15–19.
- Mengel, Friederike, and Pol Campos Mercade. 2021. “Irrational Statistical Discrimination.” *Available at SSRN 3843579*.
- Stashko, Allison. 2020. “Do Police Maximize Arrests or Minimize Crime? Evidence from Racial Profiling in U.S. Cities.” Manuscript.
- The Sentencing Project. 2015. “Black Lives Matter: Eliminating Racial Inequity in the Criminal Justice System.” <https://www.sentencingproject.org/wp-content/uploads/2015/11/Black-Lives-Matter.pdf>.