*Political Science Research and Methods* (2025), 1–18
doi:10.1017/psrm.2025.21

ORIGINAL ARTICLE

# Social segregation, inter-group contact, and discriminatory policing

Ryan Hübert[1] and Andrew T. Little[2]

[1]Department of Methodology, London School of Economics and Political Science, London, UK and [2]Department of Political Science, University of California, Berkeley, CA, USA
**Corresponding author:** Andrew T. Little; Email: andrew.little@berkeley.edu.

Authors are listed in alphabetical order. Many thanks to Bocar Ba, Allison Harris, Liqun Liu, Hannah Simpson, Jeff Staton, and Alex Stephenson for comments.

(Received 8 February 2024; revised 6 June 2024; accepted 7 August 2024)

**Abstract**
We analyze a formal model of social contact and discrimination in the context of policing. Officers decide how to interact with members of two social groups while working and while socializing. The officers do not fully distinguish between their experiences of crime across these two contexts ("coarse thinking"), so they end up with excessively positive views of groups they socialize with and excessively negative views of those they police. This creates dual feedback loops as officers choose to socialize more with groups they view favorably and over-police those they view as "more criminal." Interventions that induce positive contact with an overpoliced group can mitigate the officer's discriminatory policing. However, this beneficial effect only persists if the policy intervention creates *sustained* positive contact. Our results provide a novel theoretical microfoundation for the contact hypothesis and highlight why effects of many policy interventions aimed at increasing positive contact may be short-lived.

**Keywords:** contact hypothesis; discrimination; misspecified models; policing

In recent years, there has been renewed attention to the substantial racial disparities in a variety of domains such as policing, political campaigns, and representation. In the context of policing (particularly in the United States), many reforms have been proposed and sometimes enacted to mitigate discrimination. For example, departments could hire from underrepresented groups (Ba *et al.*, 2021), institute residency requirements for officers (Smith, 1980; Payson and Parinandi, 2025), or encourage policing that develops positive relationships with policed communities (McCall, 2019; Peyton *et al.*, 2019).

A common thread to these reforms is that they (at least in part) attempt to influence the kinds of contact that officers have with different groups. This focus is understandable, as the idea that encouraging positive contact between members of different communities may decrease prejudice and discrimination has been studied widely across the social sciences (see Pettigrew and Tropp, 2006, for an influential overview). While the reported results from these kinds of interventions are often positive, when focusing only on randomized interventions, the effects of contact are more modest (Paluck *et al.*, 2019).

From a "rationalist" perspective, even small average effects of contact on beliefs may be surprising. If we conceive of contact as giving more information about the traits of outgroups, then it will only

lead to more positive beliefs when the content of the contact is more positive than individuals expect. For this reason, under the standard assumption that prior beliefs are well calibrated, giving more information should not systematically make average beliefs more or less positive. Even if interventions are specifically designed to induce positive contact, a fully rational individual should recognize this fact and adjust accordingly.

We provide a formal model which explores how different kinds of contact on and off the job *can* affect police officers' beliefs and behavior with a mild deviation from the rationalist approach. Before getting to contact interventions, we study how the "endogenous contact" driven by the social segregation of police officers, combined with a cognitive phenomenon known as "coarse thinking" (Fryer and Jackson, 2008; Mullainathan *et al.*, 2008), can cause prejudice and discrimination in the first place.

In our context, coarse thinking means that officers do not fully account for the fact that crime rates vary across different environments. The fact that people think categorically in general is a well-established result from psychology (Macrae and Bodenhausen, 2001; Mullainathan *et al.*, 2008). The driving force behind our model is that the salient category in officers' minds is the race of the person they are interacting with rather than the context (i.e., work versus social life) in which the interaction occurs.

To be concrete, suppose an officer experiences less crime in his social life than when working and socializes primarily with members of one group. The officer will have more positive contact with that group since he generally experiences less crime when he is among members of that group than he does when he is among members of the other group whom he interacts with primarily in a policing context.

By itself this need not lead to incorrect beliefs, and for a fully Bayesian officer in our model, it will not. However, a coarse thinker who does not perfectly distinguish between different experiences in different contexts will end up with biased beliefs that are influenced by how he allocates his own time on and off the clock. Further, if we allow for officers to choose how to allocate their time based on their beliefs, this will generate spillovers between work and nonwork lives. Social segregation in an officer's personal life (i.e., spending most of his social time with members of one group) influences how he makes decisions as an officer, and his policing decisions affect how he socializes with others in his personal life. A fully rational officer correctly evaluates his past performance catching crimes and adjusts future policing behavior accordingly, eliminating any potential bias in his policing. However, a coarse thinker's incorrect inferences about crime rates prevent him from correcting potential biases in his policing and can actually exacerbate discrimination.

After showing how coarse thinking leads to discrimination in the first place, we show how it also creates an opening for interventions aimed at improving positive contact to affect beliefs. In other words, the same inferential mistakes which can cause incorrect negative beliefs about a group also provide a path for interventions to reduce prejudice. Our analysis also highlights the limitations to this approach, particularly if short-lived.

## 1. Policing disparities, incorrect beliefs, and the contact hypothesis

Racial disparities in policing have been well documented, and many statistical estimates are likely understated (e.g., Goel *et al.*, 2016; Knox *et al.*, 2020). There is a robust literature devoted to cataloging and teasing out the causes of racial disparities in policing, as well as in other domains like labor markets and politics (for example, see Anwar and Fang, 2006; Persico, 2009; Butler and Broockman, 2011; Doleac and Stein, 2013; Broockman and Soltas, 2020; Harris *et al.*, 2020; Ba *et al.*, 2021; Nathan and White, 2021).

We contribute most directly to a deep literature that examines how discrimination emerges as a by-product of individual choices and beliefs. Closest to our approach, one individual-level explanation for disparities emerges from the fact that group identity may be informative about crime. This is the mechanism driving models of statistical discrimination that have emerged from the original work of Phelps (1972) and Arrow 1973. According to these models, if crime rates are different across groups,

then certain kinds of policing objectives (like reducing crime) provide a policy rationale for policing communities differently. And to the extent that statistical discrimination is a driver of policing disparities, then policy responses aimed at reducing disparities might focus on, for example, excessive criminalization of certain behaviors.

Of course, this is not the only cause of discrimination identified in prior research. Discriminatory policing can arise if police officers are racist or otherwise prejudiced, and thus have a "taste" for discriminating (e.g., Glaser, 2015; Spriggs, 2020). In addition, Small and Pager (2020) summarize a tradition from sociology that de-emphasizes individuals' prejudices, biases, and beliefs and instead examines how discrimination "may be caused by organizational rules or by people following the law" (p. 52).

Our analysis does not preclude the possibility that discrimination may be caused by several factors. In fact, other causes of discrimination may interact with—and exacerbate—the specific cause of discrimination that we study.[1] However, a striking feature of our analysis is that a previously understudied form of belief-driven discrimination can emerge even in the absence of individual prejudice, different crime rates, or institutional factors.

A crucial open question for explanations for disparities focused on statistical discrimination is whether decision-makers' beliefs about crime rates actually match reality. Bohren *et al.* (2019) point out that it is uncommon for studies of statistical discrimination to consider (let alone test for) the possibility that discrimination could occur due to biased or otherwise incorrect beliefs. In the context of policing specifically, substantial empirical evidence suggests that some decision-makers might not have accurate beliefs about the key facts that drive policing allocations (see, for example, Glaser, 2015).

A growing literature demonstrates that inaccurate beliefs have implications for the study of discrimination. For example, Hull (2021) and Hübert and Little (2023) highlight the challenge of distinguishing between different causes of discrimination, particularly when incorrect beliefs are possible. Closest to the present paper, Hübert and Little (2023) use a similar baseline model of policing to study how a different inferential mistake—high-level police officials not accounting for the fact that crime data reflect how much effort is spent policing different communities—creates a link between taste-based discrimination and statistical discrimination, which existing work primarily treats as competing explanations.

In contrast, we focus on how categorical thinking can create a microfoundation for the contact hypothesis. In particular, we study how positive and negative contact with members from different social groups can bias police officers' beliefs about crime and cause discriminatory policing. We, therefore, connect the newly influential literature on discrimination driven by inaccurate beliefs to another influential (and mostly empirical) literature on whether positive contact with an outgroup may lead to more favorable beliefs and attitudes (for reviews, seePettigrew and Tropp, 2006; Paluck *et al.*, 2019).[2]

While our broad goal is to connect coarse thinking to contact and discrimination more broadly, policing in the United States is a ripe context to build this connection. First, not only are police officers disproportionally White and Republican, so are their neighbors (Ba *et al.*, 2022). In our theory, nonwork social segregation will play an indirect role in an officer's policing decisions by affecting how much positive contact they have with members of different social groups. Second, the conditions for contact to reduce prejudice, as initially articulated by Allport (1954)—and which have received relatively little empirical examination (seePaluck *et al.*, 2019)—are unlikely to be met in police-civilian

[1]For a more general call to study the role of race in policy feedback see Michener (2019).

[2]No studies of police met the criteria laid out by Paluck *et al.* (2019) (randomized treatment and outcomes measured at least a day afterwards), though see Samii (2013) for an example with a natural experiment on security forces in Burundi, and Peyton *et al.* (2019) for an example of how positive contact with police can affect the beliefs of citizens. In other criminal justice applications, Lerman and Mooney 2022 find that citizens coming into more contact with other racial groups may cause an increase in calls to police, and Harris (2024) finds that White judges who work in close proximity with Black judges become less punitive toward Black defendants.

interactions but are likely to be met in officers' social lives. How these two kinds of contact interact and influence officers' decision-making is the main focus of our theory.[3]

The key contribution of our model to the contact hypothesis literature is to show that attitudes polarize further when actors have the opportunity to select into positive interactions with groups they have favorable views toward while selecting into negative interactions with those they initially dislike. This "endogenous contact hypothesis" underscores the limits of interventions that do not fundamentally (and permanently) change officers' incentives and offers another avenue for empirical exploration, given that there are relatively few high quality empirical studies examine both the duration of positive contact *and* the persistence of effects over time (Paluck *et al.*, 2019).

## 2. A model of policing and socializing

We formally model the decisions of a police officer (pronoun "he"). He spends fraction $w \in (0, 1)$ of his time working on a police beat. He spends the remaining $1-w$ of his time on other activities, which we refer to interchangeably as "socializing" or "leisure." We assume that $w$ is exogenous, reflecting, for example, that he is required to work a fixed number of hours.

In both contexts, the officer comes into contact with members of two groups, $A$ and $B$. We often use $J$ to refer to a generic group. These two groups are distinguishable to the officer and could represent, for example, racial or ethnic groups, or social classes.

In reality, a police officer could spend some of his work or leisure time with groups other than $A$ and $B$ or with no one at all. However, all of our results below focus on *relative* differences between how he allocates time among members of groups $A$ and $B$. So, it does not add much to the model—beyond additional notation—to allow for contact with additional groups or other ways the officer could spend his time (e.g., doing paperwork or remaining idle). When we say "time working" we mean the time he spends working and interacting with these two groups and when saying "leisure time" we mean nonwork time spent interacting with members of these two groups.

### 2.1. Time allocation

The officer makes two choices that affect how often he interacts with members of these groups in both contexts.

First, he chooses to spend $s_A \in [0, 1]$ of his leisure time with members of group $A$ and $s_B \in [0, 1]$ of his leisure time with members of group $B$. For example, where officers choose to live is a main driver of incidental and intentional social contact, and officers in the United States do live in much Whiter and more Republican areas (Ba *et al.*, 2022). Since our model focuses on the time he spends with one of the two groups, $s_B = 1 - s_A$. Given that the officer spends $1 - w$ of his time in leisure and only interacts with these two groups, he spends $(1 - w)s_A$ and $(1 - w)(1 - s_A)$ socializing with members of groups $A$ and $B$, respectively.

Next, the officer spends $w$ of his time at work. There, he has to decide where to allocate his time and energy, which ultimately affects how much time he spends policing different groups.

We assume the officer allocates $p_A$ of his time policing group $A$ and $p_B = 1-p_A$ of his time policing group $B$.[4] To represent potential institutional and/or legal constraints on his choices, let $p_A \in [\underline{p}, \overline{p}]$, for some $0 < \underline{p} < \frac{1}{2} < \overline{p} < 1$. This formal condition gives the officer leeway to allocate his time while ensuring that he always has the possibility of choosing to police both groups equally. On the other hand, the condition rules out the possibility that he chooses to police only one of the two groups.[5]

[3]Lowe (2021) presents the results of an experiment that are largely consistent with our finding that positive interactions with outgroup members improve attitudes.

[4]This codifies the assumption at all of the work time is spent policing one group or the other.

[5]Setting bounds on the policing allocation also means the officer has some contact with both groups, which will ensure the beliefs derived in the next section are well defined.

Given that the officer spends $w$ of his total time working, he spends $wp_A$ and $w(1-p_A)$ of his total time policing members of group $A$ and $B$, respectively.

### 2.2. Crime

In the model, the incidence of crime varies on two dimensions, by context/location and by group. We abstract away from differences in types of crime and assume that all crime is uniform. For simplicity, we will simply refer to "crime," but substantively, it is easiest to think of our model as pertaining to policing of a specific crime—e.g., robbery, larceny, assault, drug trafficking, illegal possession of weapons, etc.

First, we let $\kappa_A^p \in (0,1]$ and $\kappa_B^p \in (0,1]$ represent the crime rates among members of groups $A$ and $B$ in the context the officer experiences when policing. Second, we let $\kappa_A^s \in [0,1]$ and $\kappa_B^s \in [0,1]$ be the crime rates among members of groups $A$ and $B$ as experienced in the officer's leisure time. Note the crime rates in the social context may be set to zero. That is, our results to not require officers to experience crime in their social life, though we leave open this possibility.

We primarily interpret these crime rates as probabilities, capturing how often any given interaction with a member of group $J$ will entail an experience of crime. For example, a police officer doing traffic enforcement may recover illegal contraband with some probability that varies by group.

A key assumption in our analysis is that the within-group crime rates are higher in the policing context: $\kappa_A^p > \kappa_A^s$ and $\kappa_B^p > \kappa_B^s$. This assumption captures the idea that a police department will focus its resources on places and situations where crime is more prevalent.[6] Our general analysis allows for the groups to have different crime rates within each context, which would lead an officer with correct beliefs to "statistically discriminate" in both his work and social life. To keep the main focus on incorrect beliefs as a important driver of discrimination, our main examples and illustrations will focus on the special case where the within-context crime rates are equal across the two groups (i.e., $\kappa_A^p = \kappa_B^p$ and $\kappa_A^s = \kappa_B^s$).

The officer's experiences of crime depend on the choices he makes about how to allocate his time. The amount of crime he experiences among group $J$ while at work is $c_J \equiv w\kappa_J^p p_J$. By assuming the amount of crime experienced is linear in the time spent policing the group, we abstract away from the fact that policing a group more may lower the average amount of crime per unit of time.[7]

We assume the officer utility is increasing in $c_J$ for both $J$, meaning that he aims to catch crimes committed by members of either group. This captures the notion that the officer gets utility from solving open cases, detaining offenders, etc. (Stashko, 2023). While we do not explicitly model wages or professional advancement, these also provide incentives to allocate policing time to areas with more "action." Importantly, the officer's utility from policing group $J$ is increasing in $\kappa_J^p$, and this will make the officer prefer to spend more time policing the group he believes to have a higher crime rate. Some officers may aim to "avoid action" and police areas and groups with less crime, but at a minimum our model describes the dynamics for the more aggressive individuals.[8]

On the other hand, we assume the officer prefers not to experience crime in his social life. For example, experiencing crime in his social life may make him feel anxious or cause him to feel cross-pressured between his duty as a police officer and his desire to spend time with his friends.

[6]Of course, some officers may be corrupt or otherwise engaged in criminal activity on the side, but we do not wish to center outlier cases in a generic analysis of policing.

[7]For example, if policing a group heavily means pulling them over with weaker evidence, this may lower the "hit rate" (e.g., Rambachan and Roth, 2019). Further, while we focus on the behavior of one officer, more policing of one group in the aggregate may have a deterrent effect of decreasing how often they commit crimes (Chalfin and McCrary, 2018). See Hübert and Little (2023, Online Appendices C and D) for a discussion of how these dynamic can attenuate but not eliminate the effect of a related belief bias.

[8]For example, Ba *et al.* (2021) find that officers in Chicago prefer to work in districts with lower crime rates, but transfers between districts are generally rare. Even so, our model is best understood as examining an officers incentives *conditional on* being assigned to a particular district.

Alternatively, socializing around criminal activity might cause him to lose his job or be arrested, which imposes a concrete cost with some probability that is proportional to the amount of time he spends in such a context.

More generally, we aim to formalize an idea closely related to the contact hypothesis, albeit with an additional causal arrow in the opposite direction: officers want to socialize more with groups that they expect to have positive experiences with, and less with groups that they associate with negative interactions.

To capture these ideas, we assume the officer only enjoys leisure time with individuals from group $J$ when they are not engaged in criminal activity. Formally, the officer's "positive leisure time" with group $J$ is proportional to $1 - \kappa_J^s$. Given he spends fraction $(1 - w)s_J$ of his time socializing with group $J$, the overall positive leisure time spent with members of group $J$ is $l_J \equiv (1 - w)(1 - \kappa_J^s)s_J$.

### 2.3. Utility from time spent policing and socializing

We now formally define the officer's utility function, which embeds several assumptions. Summarizing the previous section, our first core assumption about the officer's preferences is that his utility is strictly increasing in $c_J$ and $l_J$.

Second, we assume that the officer gets diminishing marginal utility from spending time on each of the four different ways he can spend his time (i.e., policing/socializing with group $A/B$). That is, his utility is strictly concave in $c_A$, $c_B$, $l_A$, and $l_B$. While the marginal utility is scaled by crime rates (recall the definitions of $c_J$ and $l_J$), we assume it always diminishes. In the case of crime, this could represent the idea that the officer detects more important crimes first, and so increasing the intensity of policing one group will lead to diminishing returns.

Third, we assume that the officer's utility is additively separable in the four different ways he can spend his time. This assumption essentially means that, with full information, more policing or socializing with one group does not affect the marginal return to the other activity or with the other group. Though, importantly, such spillovers will arise when we allow for incorrect beliefs.

Finally, we include parameters $t_A$ and $t_B$ which scale how much the officer enjoys socializing with members of groups $A$ and $B$. These capture the idea that the officer may have some group-based favoritism that affects how he spends his personal time outside of work. While we do not explicitly model which group the officer is a member of (and our analysis only focuses on one officer), differences in social circles may be an important driver of differences in officer behavior by racial group (Donohue and Levitt, 2001; Antonovics and Knight, 2009; Ba *et al.*, 2021). As a result, one could imagine that the $t_A$ and $t_B$ parameters vary across officers, including by relevant characteristics such as whether an officer is a member of group $A$ or group $B$.

It is straightforward to add similar parameters to capture the fact that the officer might prefer catching crimes committed by members of one group. For example, the officer could be racist against one group and accordingly wish to be punitive toward members of that group. While this is potentially an important driver of discriminatory policing (even when officers form incorrect beliefs, see Hübert and Little, 2023), in this paper, we keep our focus on how biases arise from the officer's social life. So that we can show that at least some portion of discriminatory policing arises due to this mechanism, we will assume that the officer does not explicitly care whether he catches crimes committed by members of one group or the other when he is at work as a police officer.

To capture these assumptions, we use the following functional form:

$$\begin{aligned} u(p_A, p_B, s_A, s_B) &= \sqrt{c_A} + \sqrt{c_B} + \sqrt{t_A l_A} + \sqrt{t_B l_B} \\ &= \sqrt{\kappa_A^p w p_A} + \sqrt{\kappa_B^p w p_B} \\ &\quad + \sqrt{t_A(1 - \kappa_A^s)(1 - w)s_A} + \sqrt{t_B(1 - \kappa_B^s)(1 - w)s_B}. \end{aligned}$$

The only consequential assumption added by this utility function is the specific use of square roots to capture the fact that the utility is increasing with diminishing returns in each component. Square roots allow us to formalize our diminishing marginal utility assumption in a way that yields tidy technical results.

### 2.4. Full information policing and socializing

We first consider how the police officer allocates his work time if he knows the (true) crime rates in the communities he polices (i.e., $\kappa_A^p$ and $\kappa_B^p$). We refer to this as "full information policing." Assuming it is an interior solution (i.e., $p_A \in (\underline{p}, \overline{p})$, the allocation of policing $p_A$ that maximizes the officer's utility is

$$p_A = \frac{r_p}{1 + r_p} \equiv p_A^{\text{br}}(r_p) \tag{1}$$

where $r_p = \kappa_A^p / \kappa_B^p$. For much of our analysis, we can summarize the crime rate parameters with this ratio capturing the relative crime rates. It is immediate that $p_A^{\text{br}}$ is increasing in $r_p$, which means the officer will want to allocate more policing time toward the group with a higher crime rate.

One important feature of this optimal allocation is that it does not depend on any of the "social life" parameters ($t_J$, $\kappa_J^s$) or choice ($s_J$). In our analysis below, the independence of his policing and socializing choices will break down when the officer does not fully distinguish between these domains when forming beliefs.

Since $r_p/(1 + r_p)$ is a number between zero and one, it follows that there is a unique policing allocation $p_A$ that solves (1). To reduce cases to consider, we assume that this lies in the officer choice set: $\underline{p} < p_A^{\text{br}}(r_p) < \overline{p}$.

The police officer also chooses how to allocate his leisure time. His maximization problem yields a similar first-order condition:

$$s_A = \frac{r_t r_s}{1 + r_t r_s} \equiv s_A^{\text{br}}(r_s, r_t) \tag{2}$$

where $r_s \equiv (1 - \kappa_A^s)/(1 - \kappa_B^s)$ and $r_t \equiv t_A/t_B$. This again yields a unique allocation $s_A^{\text{br}}$ for all $r_s$ and $r_t$. This is increasing in $r_s$, which is the ratio of "noncrime rates" of the two groups. This captures the idea that the officer prefers to spend more time socializing with members of the group he believes to have a lower crime rate. It is also increasing in $r_t$, capturing the natural effect that if the officer derives relatively more value from socializing with group $A$ he will choose to socialize more with that group.

The allocations $p_A^{\text{br}}(r_p)$ and $s_A^{\text{br}}(r_s, r_t)$ comprise the **full information benchmark**, and we label these quantities more concisely as $p_A^\dagger$ and $s_A^\dagger$. A key feature of these allocations is that they are independent optimization problems, where the officer's social choices do not affect his policing choices, and vice versa.

**Proposition 1.** *If the officer has correct beliefs about the relevant crime rates ($r_p$ and $r_s$), then:*

(i) *there is a unique optimal allocation given by $p_A^\dagger$ and $s_A^\dagger$, defined by (1) and (2), respectively.*
(ii) *The officer's policing decisions and social decisions are independent from the parameters affecting the "other" domain. Formally, $\frac{\partial p_A^\dagger}{\partial r_s} = 0, \frac{\partial p_A^\dagger}{\partial r_t} = 0$ and $\frac{\partial s_A^\dagger}{\partial r_p} = 0$.*
(iii) *There is a policing disparity ($p_A^\dagger \neq 1/2$) if and only if crime rates in the policing context are unequal ($r_p \neq 1$).*

All proofs for formal results are in Online Appendix A.

## 3. How contact can generate inaccurate beliefs about groups

So far, we have focused on an officer's policing and socializing choices, assuming that he has "full information" (and thus correct beliefs) about $r_p$ and $r_s$. But where do his beliefs come from?

### 3.1. How the officer could form correct beliefs

First consider how the officer could, in principle, correctly learn $r_p$ and $r_s$ from his experience. Learning about $r_p$ requires knowing $\kappa_A^p$ and $\kappa_B^p$, and to estimate these, the officer could ask "in what proportion of my time when policing group $J$ do I detect crime?" Formally, he could compute a belief $\tilde{\kappa}_J^p$ as follows:

$$\tilde{\kappa}_J^p(\xi = 0) = \frac{wp_J\kappa_J^p}{wp_J} = \kappa_J^p. \tag{3}$$

The numerator of equation (3) is the crime detected while policing group $J$, and the denominator is the time spent policing group $J$. That is, rather than specifying a prior belief and computing updated beliefs about crime rates via Bayes' rule, we adopt a more "frequentist" style of analysis, though under general assumptions a Bayesian would converge to this belief with a large sample (Blackwell and Dubins, 1962).[9]

The $\xi = 0$ argument is to indicate the officer forms this belief in an unbiased manner. The officer could similarly estimate the crime rate in social contexts, $\tilde{\kappa}_J^s$. And given that officer interprets his experiences of crime in an unbiased manner, neither $\tilde{\kappa}_J^p$ nor $\tilde{\kappa}_J^s$ should affect one another.

### 3.2. Coarse thinking

This logic quickly unravels when we allow for the possibility of an intuitive bias in the officer's belief formation process. Specifically, we study a situation in which the officer is a "coarse thinker" (Mullainathan, 2002) with respect to race, meaning that he (at least partially) conflates his experience of crime among group $J$ at work with his experience of crime among this group in his social life.

In the extreme case in which he is a *fully* coarse thinker (formally, $\xi = 1$), the officer estimates a single crime rate for each group, which is now endogenous to his own policing and socializing choices:

$$\underbrace{\tilde{\kappa}_J^s(\xi = 1)}_{\substack{J\text{'s crime rate in} \\ \text{his social life}}} = \underbrace{\tilde{\kappa}_J^p(\xi = 1)}_{\substack{J\text{'s crime rate} \\ \text{while policing}}} = \frac{wp_J\kappa_J^p + (1-w)s_J\kappa_J^s}{wp_J + (1-w)s_J}.$$

The numerator corresponds to the crimes experienced among group $J$ both while policing and in his social life, and the denominator is the total time spend with group $J$ in both contexts.

A more general version of these biased beliefs allows the parameter $\xi \in [0, 1]$ to scale the degree of coarse thinking, where the officer's beliefs become "coarser" as $\xi$ increases:

$$\tilde{\kappa}_J^p(\xi) = \frac{\kappa_J^p wp_J + \xi\kappa_J^s(1-w)s_J}{wp_J + \xi(1-w)s_J} \qquad \tilde{\kappa}_J^s(\xi) = \frac{(1-w)s_J\kappa_J^s + \xi wp_J\kappa_J^p}{(1-w)s_J + \xi wp_J}. \tag{4}$$

We illustrate how coarse thinking affects his beliefs in Figure 1. In all panels, the horizontal lines represent the true crime rates of group $J$ in the policing and social contexts, i.e., $\kappa_J^p$ and $\kappa_J^s$. The curves plot the officer's beliefs about these quantities, given he is a coarse thinker, i.e., $\tilde{\kappa}_J^p$ and $\tilde{\kappa}_J^s$. The left panel illustrates how the officer's beliefs about group $J$'s crime rates converge as $\xi$ increases. The right two

[9] One plausible effect this approach elides is that interacting with a group less (whether policing or socializing) could decrease the precision of the information contained about the group, leading to more volatile beliefs. Even if so, the effects we describe plausibly hold on average.

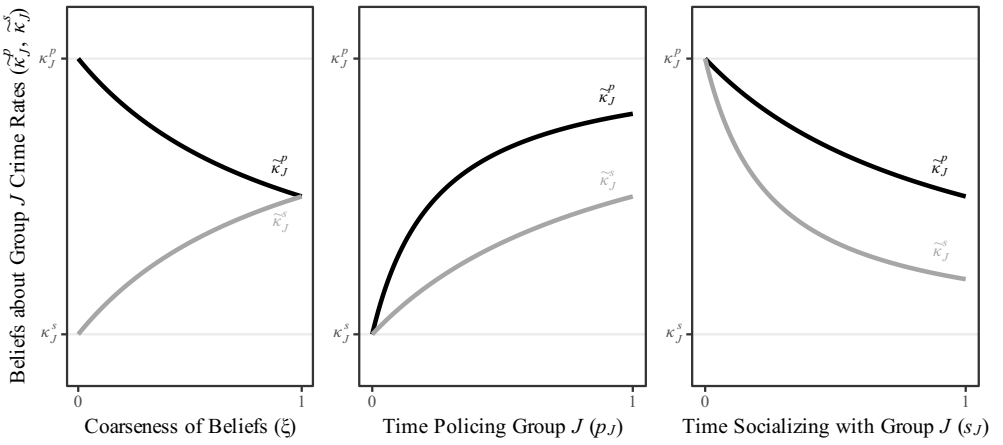


**Figure 1.** Beliefs about the crime rate of group *J* in the policing context (black) and social context (gray), as a function of the severity of his coarse thinking (left panel) policing allocation toward this group (middle panel), and social allocation (right panel).

panels illustrate how a coarse thinking officer's beliefs are (incorrectly) affected by his own policing and social decisions.

All three panels demonstrate that the officer's beliefs are almost always incorrect. More importantly, the middle panel shows that he forms a more negative view of group *J* (i.e., that the group has a higher crime rate) as he polices that group more. And the right panel shows that he forms a more positive view of group *J* as he socializes with that group more, even though the underlying crime rates don't change as he changes who he socializes with.

As in the full information case we analyzed above, what matters for our remaining analysis are the ratios of (non)crime rates in each context: $\tilde{r}_p = \tilde{\kappa}_A^p(\xi)/\tilde{\kappa}_B^p(\xi)$ and $\tilde{r}_s = (1-\tilde{\kappa}_A^s(\xi))/(1-\tilde{\kappa}_B^s(\xi))$.[10] These ratios have several interesting properties, which we show formally:

**Proposition 2.** *The officer's beliefs about the relative crime rates $\tilde{r}_p$ and $\tilde{r}_s$ have the following properties:*

*(i) If $\xi = 0$, then the officer's belief about the relative crime rates are correct: $\tilde{r}_p = r_p$ and $\tilde{r}_s = r_s$.*
*(ii) If $\xi > 0$, then the officer's beliefs about the relative crime rate in the policing context ($\tilde{r}_p$) is increasing in $\kappa_A^p$ and $\kappa_B^s$ and decreasing in $\kappa_B^p$ and $\kappa_A^s$.*
*(iii) If $\xi > 0$, then the officer's beliefs about the relative noncrime rate in the social context ($\tilde{r}_s$) is decreasing in $\kappa_A^p$ and $\kappa_B^s$ and increasing in $\kappa_B^p$ and $\kappa_A^s$.*
*(iv) If $\xi > 0$, then the officer's beliefs at work and in leisure respond to how he allocates his time in both contexts. Specifically:*

$$\frac{\partial \tilde{r}_p}{\partial p_A} > 0 \qquad \frac{\partial \tilde{r}_s}{\partial p_A} < 0 \qquad \frac{\partial \tilde{r}_p}{\partial s_A} < 0 \qquad \frac{\partial \tilde{r}_s}{\partial s_A} > 0.$$

When the officer is not a coarse thinker, his belief about the relative crime rate in the policing context is only a function of $\kappa_A^p$ and $\kappa_B^p$ (which is correct). Similarly, in the social context, his belief only depends on $\kappa_A^s$ and $\kappa_B^s$. However, when $\xi > 0$, the officer's beliefs still respond to these parameters in the expected and correct direction, but they are also sensitive to several other things that should

[10] More specifically, because the officer's utility function is homogeneous of positive degree, we are able to express the best response in terms of these ratios, see equations (1) and (2).

not matter. As stated in part (ii), increasing the crime rate of a group in the *social* context increases the officer's beliefs about that group in the policing context, and vice versa.

More important for what follows, as illustrated in Figure 1 and formalized in part (iii), the officer's beliefs become a function of both decisions: spending more time policing a group increases beliefs about their relative crime rate in both contexts and spending more time socializing with a group decreases this belief in both contexts.

### 3.3. Why this bias?

Once the Pandora's box of nonrational beliefs is opened, a natural question is which bias to focus on. We do not contend that coarse thinking is the only relevant belief for studying discrimination. Instead we contend that it is worth studying in this context as (1) there is substantial evidence for the importance of categories on belief formation, and (2) it provides a natural link between recent empirical work on policing and the contact hypothesis.

The ideas that humans think in categories and that the categories they choose influence resulting beliefs are well accepted in psychology (see Macrae and Bodenhausen, 2001, for an overview). Our model goes beyond this, however, in assuming that officers think about race when categorizing interactions but not necessarily about work/social life context. Focusing on the category of race is a reasonable assumption given its prominence in stereotypes about crime (O'Flaherty and Sethi, 2019).

Still, to be more precise, the coarseness parameter we study is really "coarseness with respect to race," rather than with respect to any potentially relevant categories. As shown in Online Appendix C, similar results arise if officers are also "coarse with respect to context," meaning they conflate experiences with group $A$ in a given context with experiences with group $B$ in the same context.

## 4. How distorted beliefs affect policing and social segregation

We have shown how coarse thinking affects the officer's beliefs about each group's crime rates. We now consider how this manifests in his behavior.

### 4.1. How social segregation causes discriminatory policing

In this section, we focus on the officer's policing behavior, considering how his social life can influence his policing behaviors when he is a coarse thinker. For now, we will treat the the officer's social life (i.e., $s_A$) as exogenous. One might consider the analysis in this section a good approximation for the "short run": how does the officer make his day to day policing decisions when he has limited ability to quickly change his social circles? In the next section, we will again allow the officer to choose how he socializes, which enables us to examine how his job as a police officer affects his social life.

Formally, in this section, we define an equilibrium as follows:

**Definition 1.** An equilibrium with fixed social life is a policing allocation $p_A^*$ and belief about the relative crime rate in the policing context $\tilde{r}_p^*$ such that:

(1) $p_A^*$ is optimal for the officer given $\tilde{r}_p^*$, and
(2) $\tilde{r}_p^*$ is calculated using the left equation in (4) with policing allocation $p_A^*$.

To further simplify and generate a relatively straightforward closed-form solution, we also impose a restriction that $\kappa_J^s = 0$, i.e., the real rate of crime for both groups in the social context is zero. Under these assumptions, there is a unique potential interior allocation which solves $p_A^I = p_A^{\text{br}}(\tilde{r}_p(p_A^I))$,

given by:

$$p_A^I = \underbrace{\frac{r_p}{1+r_p}}_{p_A^\dagger} + \underbrace{\xi\left[\frac{1-w}{w}\right]\left[\frac{r_p(1-s_A)-s_A}{1+r_p}\right]}_{\text{bias term}}. \tag{5}$$

This policing allocation is equal to the full information benchmark $p_A^\dagger$ plus a bias term that can be either positive or negative. If it is positive (negative), then he polices group $A$ (group $B$) more than he would do under full information. Moreover, if he is a coarse thinker ($\xi > 0$), the bias term is always nonzero except for the knife-edge situation in which $s_A$ just happens to equal $p_A^\dagger$.

**Proposition 3.** *Suppose that $s_A$ is fixed and that $\kappa_j^s = 0$ for all* J. *Then, there is a unique equilibrium with a policing allocation $p_A^*$ and a belief $\tilde{r}_p^*$. At this allocation:*

*(i) If $s_A = p_A^\dagger$, then $p_A^* = p_A^\dagger$ and the officer has a accurate belief about the relative crime rates of the groups, $\tilde{r}_p^* = r_p$.*

*(ii) If $s_A < p_A^\dagger$, then $p_A^* = \min\{p_A^I, \overline{p}\} > p_A^\dagger$, and the officer overestimates the relative crime rate among members of group* A, *$\tilde{r}_p^* > r_p$.*

*(iii) If $s_A > p_A^\dagger$, then $p_A^* = \max\{p_A^I, \underline{p}\} < p_A^\dagger$, and the officer underestimates the relative crime rate among members of group* A, *$\tilde{r}_p^* < r_p$*

*(iv) Where the equilibrium is interior ($p_A^* \in (\underline{p}, \overline{p})$), $p_A^*$ is decreasing in the time socializing with group* A *and increasing in the time socializing with group* B. *Formally, $\frac{\partial p_A^*}{\partial s_A} < 0$.*

*(v) Where the equilibrium is interior ($p_A^* \in (\underline{p}, \overline{p})$), as the crime rate of group* A *in the policing context increases or the crime rate of group* B *in the policing context decreases, $p_A^*$ is increasing. Formally, $\frac{\partial p_A^*}{\partial r_p} > 0$.*

For example, for the case where crime rates are equal ($r_p = 1$), then in order for the officer to have an accurate belief about the relative crime rate and choose to police at the full information benchmark, it must be that $s_A = 1/2$. Intuitively, the officer will end up "overpolicing" the group with which he socializes less (relative to the group with which he socializes more) even though the facts on the ground—i.e., crime rates as reflected by $r_p$—don't provide justification for policing that group more heavily.

This result highlights one mechanism by which diversifying a police force may alter police behavior. If, for example, police officers from group $A$ socialize with group $A$ individuals more often than police officers from group $B$ do (i.e., have a higher $s_A$), then Proposition 3 implies that they will discriminate less against group $A$ (for empirical evidence consistent with this, see Ba *et al.*, 2021).

Figure 2 illustrates. The left panel plots the officer's best response for the case where crime rates across groups are equal ($r_p = 1$), but where the officer is a coarse thinker ($\xi = 1/2$) and socializes with group $B$ more than group $A$ ($s_A = 1/4$, $s_B = 3/4$). An equilibrium allocation occurs where the best response crosses the 45-degree line. Since $r_p = 1$, if the officer had accurate beliefs ($\xi = 0$), then he would police the groups equally: at the gray dashed lines where $p_A = 1/2$. However, given that the officer is a coarse thinker and he socializes more with group $B$, this causes him to overestimate the relative crime rate among members of group $A$, and hence the equilibrium allocation involves discrimination against this group (the black dashed lines, where $p_A > 1/2$).

While our discussion above has focused on the special case where crime rates in the policing context are equal ($r_p = 1$), our results apply even when crime rates in the policing context are unequal. For example, if $r_p < 1$, then the officer will police group $B$ more heavily all else equal. However, if the

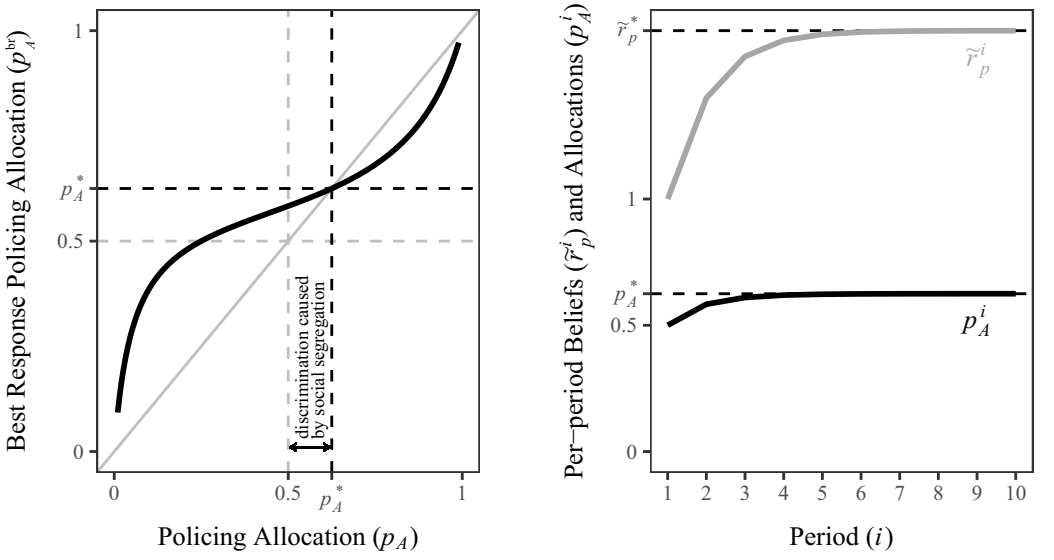


**Figure 2.** The effect of the police officer's coarse thinking on his policing allocation and beliefs about crime when he socializes primarily with members of group *B* and when crime rates across groups are equal. Specifically, we assume that $s_A = 1/4$, $s_B = 1 - s_A = 3/4$, and $r_p = 1$.

officer also socializes with group *A* more than group *B*, he polices group *B* even more heavily than he would due solely to the difference in crime rates.

To preview some later analysis and discussion, we can also think of the equilibrium beliefs and behaviors as the result of a dynamic process. For example, suppose he were to "start" by policing both groups equally. Since he socializes primarily with group *B* (recall we assume $s_A = 1/4$ and $s_B = 3/4$) and is a coarse thinker ($\xi = 1/2$), the belief generated with equal policing would make him think the crime rate of group *A* is higher, and hence induce him to decide to police group *A* more. Indeed, we can see that the black best response curve in the left panel is higher than the 45-degree line if he were to choose $p_A = 1/2$, indicating that he would prefer to police group *A* more than he currently is. The right panel of Figure 2 illustrates what happens if this process is iterated over several "periods" indexed by $i$ where the officer forms a belief about the relative crime rates in each period ($\tilde{r}_p^i$, the gray curve) based on what happened in the previous period, and then chooses his policing allocation ($p_A^i$, the black curve) accordingly. The officer's belief $\tilde{r}_p^i$ increases in each period $i$ until it eventually converges to the equilibrium in which he overpolices group *A* and (incorrectly) believes the crime rate is higher among this group.

### 4.2. How discriminatory policing causes social segregation

We have shown how segregation in the officer's social life can cause discriminatory policing. However, we also showed that coarse thinking causes convergence between the officer's beliefs about crime while policing and his beliefs about crime in his social life. This suggests that the causal arrow also goes in the other direction: his experiences policing can influence how he thinks about the desirability of socializing with the two groups.

To explore this, we now return to the model where both the policing and social allocations are endogenous choices made by the officer, which we already analyzed for the simple case where the officer is not a coarse thinker (i.e., the full information case). Formally:

**Definition 2.** An equilibrium with endogenous policing and socializing consists of a pair of allocation choices $(p_A^*, s_A^*)$ and beliefs about relative crime rates $\tilde{r}_p(p_A^*, s_A^*)$ and $\tilde{r}_s(p_A^*, s_A^*)$, such that:

(1) $(p_A^*, s_A^*)$ are optimal allocations given $\tilde{r}_p(p_A^*, s_A^*)$ and $\tilde{r}_s(p_A^*, s_A^*)$, and
(2) $\tilde{r}_p(p_A^*, s_A^*)$ and $\tilde{r}_s(p_A^*, s_A^*)$ are calculated using equation 4 with allocations $(p_A^*, s_A^*)$.

Intuitively, this definition says that the officer behaves as if the coarse beliefs derived above are the truth, not accounting for how they are influenced by his behavior. An equilibrium of this form always exists:

**Proposition 4.** *There is an equilibrium with policing and social allocations* $(p_A^*, s_A^*)$ *such that* $p_A^*$ *and* $s_A^*$ *are mutually reinforcing, given beliefs* $\tilde{r}_p$ *and* $\tilde{r}_s$.

Using the officer's best response functions (see equations 1 and 2) and substituting his belief, a solution (if it is interior on both choices) can be characterized by:

$$p_A^* = \frac{\tilde{r}_p(p_A^*, s_A^*)}{1 + \tilde{r}_p(p_A^*, s_A^*)} \qquad s_A^* = \frac{r_t \tilde{r}_s(p_A^*, s_A^*)}{1 + r_t \tilde{r}_s(p_A^*, s_A^*)}.$$

Unfortunately, there is no explicit solution to this system of equations. However, we still can describe how the key exogenous parameters—in particular, the real crime rates as captured by $r_p$ and the officer's favoritism in his social life as captured by $r_t$—affect the officer's behavior at any equilibrium where the allocations are interior (i.e., the officer does not police or socialize with one group as much as possible) and which meet a standard stability condition.[11]

**Proposition 5.** *In the model where both the policing and social allocations are endogenous, at any stable interior allocation:*

(i) *As the officer's social favoritism toward group* A *increases, he socializes with group* A *more, and if* $\xi > 0$ *he polices group* A *less.*
(ii) *As the crime rate of group* A *in the policing context increases or the crime rate of group* B *in the policing context decreases, the officer spends more time policing group* A *and less time socializing with group* A.

While there may be multiple equilibria to the model—a point we discuss in Online Appendix B—the comparative statics of Proposition 5 hold within any stable equilibrium. In words, this result means that when we reintroduce the possibility that the officer will endogenously choose how he allocates time in his social life, this magnifies the spillover between his policing and socializing choices as there are now two feedback loops. In particular, his discriminatory policing causes him to socially segregate and his social segregation causes him to police in a discriminatory manner. These choices are based on his (unjustifiably) negative views about the relative propensity of one group's members to commit crimes.

One implication of Proposition 5 relates to the behavior of officers from different social groups. For example, if police officers have some taste for in-group social interactions, then Officer 1 from group *A* will have a higher $r_t$ than Officer 2 from group *B*. Then, part (i) of the proposition implies

[11] Online Appendix A.1 contains a formal definition, but the main idea is that starting at any equilibrium allocation, an exogenous perturbation of the officer's strategy (either the policing choice or socializing choice, or both) will cause the officer to best respond in a way that moves him back toward the equilibrium allocation.

that Officer 1 will socialize more with group $A$ than Officer 2. Moreover, if $\xi > 0$, Officer 1 will also police group $A$ less than Officer 2.[12]

## 5. Policy interventions and the contact hypothesis

To the extent that coarse thinking is an empirically common phenomenon, this creates a particular challenge for policing. Almost by definition, the work of policing entails *negative* contact with other individuals. As we've shown in our analysis, if this negative contact is disproportionate across groups, this can kick off feedback loops that lead to more social segregation and more discriminatory policing.

For a policy maker interested in counteracting these feedback loops, the fact that policing routinely entails negative contact constrains the set of policy responses since she cannot realistically direct officers to have substantially fewer negative interactions while doing police work. Instead, the policy maker could try to counteract officers' negative interactions on the job with positive *social* interactions. For example, she could implement programs that incentivize officers to spend time building positive social ties with members of overpoliced groups. Alternatively, she could hire additional officers from an overpoliced group. This latter policy could create opportunities for positive social contact among officers of different backgrounds, as well as increase the share of officers whose own social lives cause them to have less biased beliefs in the first place.

Returning to our model, consider what would happen if the officer were subjected to a policy intervention that aims to increase his positive contact with group $A$. In the main text, we will again simplify the exposition by assuming the officer's social allocation ($s_A$) is exogenous. However, in Online Appendix D, we show that the basic take-aways remain when we allow the officer to endogenously choose his social allocations (albeit with more moving parts and notation).

So, suppose this intervention increases the amount he socializes with members of group $A$ ($s_A$). In this situation, Proposition 3 might give reason to expect such an intervention could work at reducing policing disparities. Since the officer interacts with group $A$ more often while on the beat, this causes him to have excessively pessimistic views about crime among this group. In principle, this could be corrected by increasing the amount of positive interactions. However, whether this persists depends heavily on whether the policy intervention that increases his positive contact with group $A$ is permanent.

In Figure 3, we use the same parameter values as in Figure 2—namely, that $r_p = 1$, $s_A = 1/4$, and $\xi = 1/2$. Each panel depicts a 10-period dynamic process in the spirit of the right panel of Figure 2 except that they start at the equilibrium allocation depicted by the dashed lines in that figure (which here we label $p_A^0$ and $\tilde{r}_p^0$). In particular, we assume here that the officer chooses an allocation for some time period, forms a new belief based on the experiences he has given that allocation, and then in the next time period chooses the optimal allocation given this new belief.

The left panel shows what happens when the officer's positive contact with group $A$ is temporarily increased for only one period. In particular, suppose that there is an intervention in which he is required to spend half of his time socializing with group $A$ during the first period (and only the first period). Since we started by assuming that he spent a quarter of his time socializing with group $A$, this effectively doubles his positive contact with group $A$ for that one period.

This temporary change to his social life leads to a sharp change in his beliefs and behavior immediately after period 1, since he comes to believe the groups' crime rates are more equal than he thought before. However, since the intervention is only temporary, the officer (exogenously) goes back to his relatively more unequal social life in the next period. As a result, he eventually returns to his previously biased beliefs and discriminatory policing. The middle panel shows what happens if the intervention lasts more than one period but is still temporary. A similar pattern as the left panel

[12]It is also possible that $\xi$ varies by officer in ways that are correlated with officers' identities.

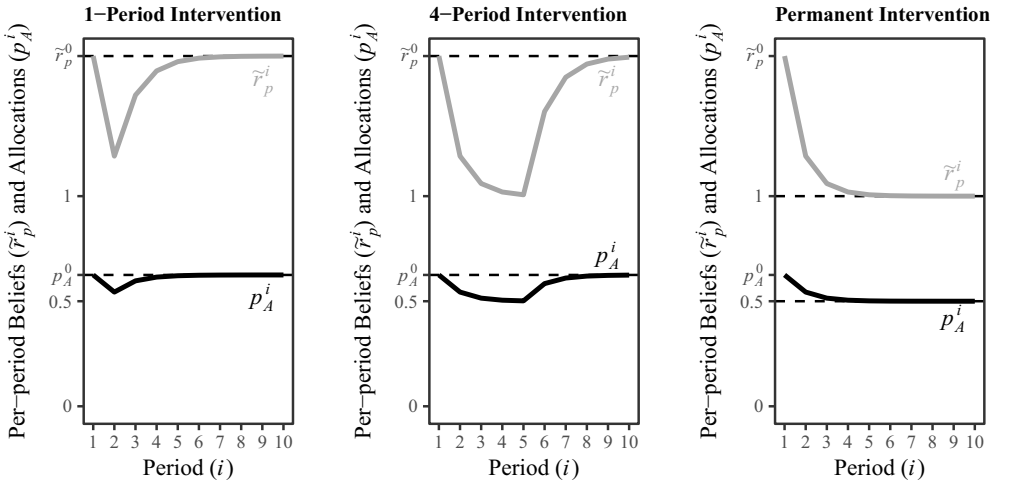


**Figure 3.** Dynamics of policing allocations and beliefs after three different policy interventions. In the left panel, the policy intervention is a very short term (i.e., 1-period) increase in how much he socializes with group *A*; in the middle panel, the policy intervention is a longer term (i.e., 4-period) increase in how much he socializes with group *A*; and in the right panel, the policy intervention is a permanent increase in how much he socializes with group *A*.

emerges, although the change in behavior caused by the intervention persists for longer since the intervention is longer.

The right panel shows what happens when the officer is subjected to a permanent change to his social life. In this case, we assume the intervention permanently moves $s_A$ from $1/4$ to $1/2$. The effect of this change actually *increases* over time, as the officer polices group $A$ less and less as his beliefs about the relative crime rates get more and more accurate. Eventually, this permanent intervention causes his behavior to converge to a new permanent equilibrium with less discriminatory policing and less social segregation.

Interestingly, interventions like these try to alleviate discrimination by taking advantage of the belief bias that caused the discrimination in the first place. If the officer were not a coarse thinker, the positive interactions caused by the policy intervention would not influence his beliefs about crime in the policing context and thus wouldn't affect how he makes policing decisions.

### 5.1. Implications for specific policies

In general, the highest quality empirical evidence on interventions to increase positive contact demonstrates modest reductions in discriminatory attitudes or behaviors (Paluck *et al.*, 2019). However, relatively few of these studies examine the effect of sustained contact or measure whether effects attenuate over time. Our analysis highlights an important and under-explored factor that may influence whether certain kinds of interventions have an impact on long-term behavior—namely, whether the intervention permanently changes the parameters that determine equilibrium behavior.

In the context of policing, programs which lead to positive interactions between police and communities for a single event, or for time spans on the order of days and weeks, likely have a modest impact on police officers' overall beliefs about that group. The positive contact may be memorable and salient in the short-term, and in most experimental work on the contact hypothesis, this is when outcomes are measured. But once officers go back to their "normal" interactions, our model suggests that any effect of the intervention will diminish unless it changes beliefs so dramatically that it actually shifts the officer to a different equilibrium. Of course, the latter is only a possibility if multiple equilibria exist in the first place.

In contrast, if an intervention consistently and permanently puts officers into more positive contact with groups while working (e.g., a community policing approach), this could have a corrective effect on officers' beliefs that persists, and even magnifies, over time as officers adjust their other work (and potentially social) behavior. That said, even though our model suggests long-term interventions can yield this particular benefit, we emphasize that there still may be downsides to such interventions, especially if they do not have wide-spread support within policed populations (Blair *et al.*, 2021).

A related class of policy reforms which have been heavily studied in the context of policing is increasing the diversity of officer ranks. Across a variety of contexts, officers are typically less aggressive and punitive toward in-group members (e.g., Antonovics and Knight, 2009; Ba *et al.*, 2021; Harvey and Mattia, 2024). Of course, policy interventions like this are best considered "compound treatments" that potentially affect outcomes through a variety of causal pathways. But one plausible set of pathways relates to the equilibrium dynamics we study above. For example, police officers from historically underrepresented groups may be more likely to socialize with (and hence have more positive views toward) in-group members. Efforts to hire more officers like this may therefore directly reduce discriminatory policing.

Perhaps more interestingly, hiring a more diverse police force can also affect the behavior of *other* officers for reasons specifically tied to our theory. For example, Rivera (2024) shows that police academy trainees who are randomly assigned to have more peers from minority groups in their training cohort subsequently make fewer "low-quality" arrests of members of those minority groups. Samii (2013) finds similar results in a very different context: members of the Burundian military who were exposed to a more diverse set of colleagues following an integration policy expressed less prejudice.

Future empirical work and policy interventions could test our theory more explicitly in several ways. One direction would be to vary the duration of positive contact and measuring beliefs and behavior of officers over time. We are also unaware of any work which looks at the relationship between where officers live (or, more generally, who they socialize with) and their behavior on the job.

Finally, on the more micro-level, future work could study whether those more predisposed to coarse thinking are more prone to prejudicial beliefs and also more strongly affected by policy interventions aimed at increasing positive contact. Work in this vein could also attempt to measure the beliefs that officers hold about the prevalence of crime among different groups, and to see if this correlates with where they live, their own racial group, and behavior on the job. Finally, it may be possible to design interventions to reduce coarse thinking; see Dube *et al.* (2025) for evidence that a training program designed to help officers think through alternative explanations for situations led to fewer discretionary arrests, particularly of Black subjects.

## 6. Conclusion

In this paper, we analyze a model of policing in which a police officer's social life influences how he makes decisions while policing two different communities. We relax the standard rational choice assumption that the police officer has to have *correct beliefs* about crime rates across communities. Instead, we allow the officer to have an intuitive behavioral bias in which he partially conflates his experiences of crime while on the beat with his experiences of crime in his nonwork social life.

Our analysis shows how his decisions about how he allocates his time in his social life spill over and affect how he polices members of two social groups. This occurs even though we assume that the officer's social life and work life are completely separate from one another. As a result, the officer ends up overpolicing members of the group he spends less time socializing with in his personal life *and* spends less time socializing with members of the group that he observes committing more crimes when he is working as a police officer.

Finally, while our analysis focuses on the example of policing, it clearly has applications to other contexts where discriminatory attitudes and behaviors are a source of concern and where positive

contact may mitigate discrimination. For example, bureaucrats and other workers in domains such as education (Levy and Razin, 2017), social work, and other areas of the criminal justice system may have disproportionately negative (or positive) contact with individuals due to the nature of their jobs.[13] Applying the theoretical ideas to these domains also expands the range of potential empirical tests.

**Supplementary material.** The supplementary material for this article can be found at https://doi.org/10.1017/psrm.2025.21. To obtain replication material for this article, https://doi.org/10.7910/DVN/GPST3C.

**Competing interests.** The author(s) declare none.

## References

**Allport GW** (1954) *The Nature of Prejudice*. Reading, MA: Addison-Wesley.

**Antonovics K and Knight BG** (2009) A new look at racial profiling: Evidence from the Boston Police Department. *The Review of Economics and Statistics* **91**, 163–177.

**Anwar S and Fang H** (2006) An alternative test of racial prejudice in motor vehicle searches: Theory and evidence. *The American Economic Review* **96**, 127–151.

**Arrow KJ** (1973) The theory of discrimination. In O Ashenfelter, and A Rees (eds), *Discrimination in Labor Markets*. Princeton, NJ: Princeton University Press.

**Ba B, Bayer P, Rim N, Rivera R, Sidibé M** (2021) Police officer assignment and neighborhood crime. 29243.

**Ba B, Kaplan J, Knox D, Komisarchik M, Mariman R, Mummolo J, Rivera R, and Torres M** (2022) Who are the police? Descriptive representation in the coercive arm of government https://dcknox.github.io/files/BaEtAl_WhoAreThePolice.pdf.

**Ba B, Knox D, Mummolo J and Rivera R** (2021) The role of officer race and gender in police-civilian interactions in Chicago. *Science* **371**, 696–702.

**Blackwell D and Dubins L** (1962) Merging of opinions with increasing information. *The Annals of Mathematical Statistics* **33**, 882–886.

**Blair G, Weinstein JM, Christia F, Arias E, Badran E, Blair RA, Cheema A, Farooqui A, Fetzer T, Grossman G, Haim D, Hameed Z, Hanson R, Hasanain A, Kronick D, Morse BS, Muggah R, Nadeem F, Tsai LL, Nanes M, Slough T, Ravanilla N, Shapiro JN, Silva B, Souza PCL and Wilke AM** (2021) Community policing does not build citizen trust in police or reduce crime in the Global South. *Science* **374**, 1–14.

**Bohren JA, Haggag K, Imas A, and Pope DG** (2019) Inaccurate statistical discrimination: an identification problem. NBER Working Paper No. 25935 https://www.nber.org/papers/w25935.

**Broockman DE and Soltas EJ** (2020) A natural experiment on discrimination in elections. *Journal of Public Economics* **188**, 104201.

**Butler DM and Broockman DE** (2011) Do politicians racially discriminate against constituents? A field experiment on state legislators. *American Journal of Political Science* **55**, 463–477.

**Chalfin A and McCrary J** (2018) Are U.S. cities underpoliced? Theory and evidence. *The Review of Economics and Statistics* **100**, 167–186.

**Doleac JL and Stein LC** (2013) The visible hand: Race and online market outcomes. *The Economic Journal* **123**, F469–F492.

**Donohue JJ and Levitt SD** (2001) The impact of race on policing and arrests. *The Journal of Law and Economics* **44**, 367–394.

**Dube O, MacArthur SJ and Shah AK** (2025) A cognitive view of policing. *The Quarterly Journal of Economics* **140**, 745–791.

**Fryer R, and Jackson MO** (2008) A categorical model of cognition and biased decision making. *The B.E. Journal of Theoretical Economics* **8** Article 6.

**Glaser J** (2015) *Suspect Race: Causes and Consequences of Racial Profiling*. New York: Oxford University Press.

**Goel S, Rao JM and Shroff R** (2016) Precinct or prejudice? Understanding racial disparities in New York City's stop-and-frisk policy. *The Annals of Applied Statistics* **10**, 365–394.

**Harris AP** (2024) Can racial diversity among judges affect sentencing outcomes? *American Political Science Review* **118**, 940–955.

**Harris AP, Ash E and Fagan J** (2020) Fiscal pressures and discriminatory policing: Evidence from traffic stops in Missouri. *Journal of Race, Ethnicity, and Politics* **5**, 1–31.

**Harvey A and Mattia T** (2024) Reducing racial disparities in crime victimization: Evidence from employment discrimination litigation. *Journal of Urban Economics* **141**, 1–27.

**Hübert R and Little AT** (2023) A behavioural theory of discrimination in policing. *The Economic Journal* **133**, 2828–2843.

**Hull P.** (2021) What marginal outcome tests can tell us about racially biased decision-making. 28503.

[13] Similar dynamics may also apply in labor markets (Fryer and Jackson, 2008) and even political competition (Levy *et al.*, 2022).

**Knox D, Lowe W and Mummolo J** (2020) Administrative records mask racially biased policing. *American Political Science Review* **114**, 619–637.
**Lerman AE and Mooney A** (2022) Racial threat, demographic change, and police calls for service. 1-38. https://possibilitylab.berkeley.edu/do-service-calls-to-911-perpetuate-racialized-policing/.
**Levy G and Razin R** (2017) The coevolution of segregation, polarized beliefs, and discrimination: The case of private versus state education. *American Economic Journal: Microeconomics* **9**, 141–70.
**Levy G, Razin R and Young A** (2022) Misspecifed politics and the recurrence of populism. *American Economic Review* **112**, 928–962.
**Lowe M** (2021) Types of contact: A field experiment on collaborative and adversarial caste integration. *American Economic Review* **111**, 1807–44.
**Macrae CN and Bodenhausen GV** (2001) Social cognition: Categorical person perception. *British Journal of Psychology* **92**, 239–255.
**McCall A** (2019) Resident assistance, police chief learning, and the persistence of aggressive policing tactics in black neighborhoods. *Journal of Politics* **81**, 1133–1142.
**Michener J** (2019) Policy feedback in a racialized polity. *Policy Studies Journal* **47**, 423–450.
**Mullainathan S** (2002) A memory-based model of bounded rationality. *The Quarterly Journal of Economics* **117**, 735–774.
**Mullainathan S, Schwartzstein J and Shleifer A** (2008) Coarse thinking and persuasion. *The Quarterly Journal of Economics* **123**, 577–619.
**Nathan NL, and A White** (2021) Experiments on and with street-level bureaucrats. In J Druckman, and D Green (eds), *Handbook of Advances in Experimental Political Science*, Chapter 27. Cambridge: Cambridge University Press.
**O'Flaherty B and Sethi R** (2019) *Shadows of Doubt: Stereotypes, Crime, and the Pursuit of Justice*. Cambridge, MA: Harvard University Press.
**Paluck EL, Green SA and Green DP** (2019) The contact hypothesis re-evaluated. *Behavioural Public Policy* **3**, 129–158.
**Payson J and Parinandi S** (2025) Residency blues: The unintended consequences of police residency requirements. *The Journal of Politics* **87**, 320–334.
**Persico N** (2009) Racial profiling? Detecting bias using statistical evidence. *Annual Review of Economics* **1**, 229–253.
**Pettigrew TF and Tropp LR** (2006) A meta-analytic test of intergroup contact theory. *Journal of Personality and Social Psychology* **90**, 751.
**Peyton K, Sierra-Arévalo M and Rand DG** (2019) A field experiment on community policing and police legitimacy. *PNAS* **116**, 19894–19898.
**Phelps ES** (1972) The statistical theory of racism and sexism. *American Economic Review* **62**, 659–661.
**Rambachan A, and Roth J** (2019) Bias in, bias out? Evaluating the folk wisdom. arXiv preprint arXiv:1909.08518. https://doi.org/10.48550/arXiv.1909.08518.
**Rivera R** (2024) Do peers matter in the police academy? *American Economic Journal: Applied Economics* **16**, 672–701.
**Samii C** (2013) Perils or promise of ethnic integration? Evidence from a hard case in Burundi. *American Political Science Review* **107**, 558–573.
**Small ML and Pager D** (2020) Sociological perspectives on racial discrimination. *Journal of Economic Perspectives* **34**, 49–67.
**Smith DC** (1980) Police attitudes and performance: The impact of residency. *Urban Affairs Quarterly* **15**, 317–334.
**Spriggs B** (2020) *Is Now a Teachable Moment for Economists*. Minneapolis, MN: Federal Reserve Bank of Minneapolis, Minnesota.
**Stashko A** (2023) Do police maximize arrests or minimize crime? Evidence from racial profiling in U.S. cities. *Journal of the European Economic Association* **21**, 167–214.

**Cite this article:** Hübert R and Little AT (2025) Social segregation, inter-group contact, and discriminatory policing. *Political Science Research and Methods*, 1–18. https://doi.org/10.1017/psrm.2025.21