Improving Hot Spot Policing through Behavioral Interventions

Final Summary Overview

Department of Justice, Office of Research and Evaluation, National Institute of Justice

Grant award #	2013-R2-CX-0006
Grantee Organization	University of Chicago
Grant Start Date and End Date	January 1, 2014 - December 31, 2018
Project Title	Improving Hot Spot Policing through Behavioral Interventions
Principal Investigator (PI)	Anuj Shah
PI e-mail address	anuj.shah@chicagobooth.edu
PI phone number	773-834-2342

Acknowledgements. This project was supported by Award No. 2013-R2-CX-0006, awarded by the National Institute of Justice, Office of Justice Programs, U.S. Department of Justice. The non-public NYPD data used in this project were provided by and belong to the NYPD. Any further use of these data must be approved by the NYPD. The publicly-available NYPD data used in this project may be accessed at the National Archive of Criminal Justice Data (NACJD). Points of view or opinions contained within this document are those of the author and do not necessarily represent the official position or policies of the Department of Justice nor those of the New York City Police Department.

We are grateful to the New York City Police Department, particularly Tara Coffey and Doug Williamson in the Office of Management Analysis and Planning. We would also like to thank Alexandra Alhadeff, Roseanna Ander, Aaron Chalfin, Katy Falco, Valentine Gilbert, Lucie Parker, David Hafetz, Zubin Jelveh, and Jason Lerner for their critical support in the planning and evaluation of the policing intervention. We also appreciate the insightful feedback on this project from Sendhil Mullainathan, Ed O'Brien, and Felipe Goncalves.

INTRODUCTION

A majority of large police departments report using some sort of hot spot strategy (Braga, 2008), which concentrates resources in high-crime areas. Hot spot strategies are based on the same rational-actor framework that shapes criminal justice policy more generally. This framework assumes that offenders choose to commit crimes based on careful consideration of the expected costs and benefits of the offense. This framework suggests that increasing the probability and cost of apprehension ought to deter would-be offenders. While most randomized controlled trials (RCTs) demonstrate that hot spot policing appears to reduce crime, some studies suggest the impacts could be modest, and several studies suggest the effects may decay after the additional police presence is withdrawn.

Clearly, there is room for improvement. The existing framework provides little insight into the factors that shape offenders' perceptions of the costs and benefits of a crime. And it has nothing at all to say about how offenders might consider other factors besides the costs and benefits of a crime. This project—a collaboration between the University of Chicago Crime Lab, ideas42, and NYPD—sought to address these gaps by using a behavioral science approach to understanding offender decision-making in hot spots in New York City.

We conducted this research in three phases. Phase 1 developed a set of hypotheses about offender decision-making. To begin, we conducted semi-structured interviews with individuals who were currently incarcerated, formerly incarcerated individuals, individuals currently on probation, and community members of high crime areas with no justice-involvement. These interviews suggested several factors worthy of further testing. For instance, offenders believed they were less likely to get away with a crime if they knew more about the officers in their community. That is, when police officers were less anonymous, offenders said they were less likely to go forward with a

crime. They were also less likely to commit a crime if they thought that children were in the vicinity. And they were less likely to commit a crime if they were in an unfamiliar setting. For Phase 2, we tested each of these hypotheses through a series of laboratory experiments and selected the hypothesis that seemed most promising for developing a field intervention (based on the results from the lab experiments)—the hypothesis about officer anonymity.

In Phase 3, we developed and conducted a field intervention to test whether reducing officer anonymity might deter crime. We worked with NYPD officers who work in New York City Housing Authority (NYCHA) developments. In a randomized controlled trial (RCT), we tested whether sending information about officers to residents in housing developments would deter crime in those developments. Preliminary results suggest that the intervention was indeed effective, particularly in the areas surrounding the developments.

Below, we describe our research activities in more detail. First, we discuss the hypothesis development and qualitative interviews from Phase 1. Then, we discuss hypothesis testing using lab-style experiments from Phase 2. Finally, we describe the field intervention from Phase 3.

PHASE 1: QUALITATIVE INTERVIEWS FOR HYPOTHESIS DEVELOPMENT

We conducted qualitative interviews with 28 individuals, drawn from four populations: 8 currently incarcerated individuals, 10 formerly incarcerated individuals, 6 individuals currently on probation, and 4 community members of high crime areas with no justice-involvement.[1] Interviews covered topics related to experiences with and perceptions of policing, safety and

[1] Interviewees were identified with support from community organizations in Harlem and East Harlem and the NYC Dept. of Correction. This is an admittedly small sample size due to the time-intensive nature of the interviews. The goal of these interviews was not to provide a representative or comprehensive overview of offender decision-making, but rather to identify themes and hypotheses that had the potential to be broadly applicable. The subsequent phases then tested these hypotheses more rigorously to determine whether they were worth exploring further.

criminal activity in the community, experiences with incarceration, and decision-making before, during and after an offense. The research team collected over 30 hours of interviews and qualitatively identified general themes for this phase of the study.

The most informative segment of the interviews focused on decision-making during the act of committing an offense. Justice-involved respondents described the context, features of the environment, and decision-making processes that led them to ultimately commit an offense or to change their mind and not go through with the act. The responses to these questions ultimately became the primary starting point for our hypothesis generation. We considered common, as well as surprising, themes among respondents' answers and used them to generate several candidate hypotheses.

For example, respondents often noted that when they saw or thought of children, they changed their mind about committing an offense. This led us to develop the hypothesis that interventions that reminded people of children (not necessarily one's own) might deter crime. Another hypothesis was based on the fact that some respondents mentioned that when they were in familiar settings, they felt more comfortable. This led us to hypothesize that being in a familiar context leads to overconfidence in decision-making and a false sense of security, both for offenders and also for victims.

However, the most promising hypothesis was that when police are more identifiable or individuated, they are seen as more trustworthy and it is more likely that an individual will think of the officer before committing an offense. For example, a young man incarcerated in Rikers Island said, *"In my neighborhood, I know the officers that work, there's one we know in my neighborhood called 'Birdman' we call him 'Birdman.' Whenever we're going to sell drugs or jump the train, and we go 'Oh Birdman's here, we got to chill, pay the fare.'"* Based on responses

such as this, we hypothesized that potential offenders might avoid committing a crime when they can easily think of specific officers in their community.

PHASE 2: TESTING HYPOTHESES WITH LABORATORY EXPERIMENTS

We conducted a series of online experiments to more rigorously evaluate the hypotheses described above.[2] For example, to test whether thinking of children could deter crime, we randomly assigned people to either think about how often they see young children in their neighborhood or to write about how often they pray or go to a house of worship (this condition was used as a control condition that might prime people to think about morality or ethics). We then asked people how likely they would be to engage in a number of aggressive or violent behaviors (e.g., getting into an argument with tailgating driver, throwing something at a dishonest relationship partner). These examples were selected to be examples of undesirable or moderately anti-social behaviors that were relevant to the participants. We predicted that participants who were primed to think about children would be less likely to endorse these aggressive behaviors. However, we did not find a significant difference in responses across the two conditions. In another study, we primed participants by either showing them pictures of children or houses of worship, but again did not find a significant difference. To test whether unfamiliar settings would make people feel less safe or comfortable, we showed participants scenes of different parts of a city. Some scenes were shown many times (to increase familiarity with those scenes) and some scenes were shown few times. Participants then rated how comfortable they would feel in each of the scenes. We expected that

[2] These experiments were intended to test whether the hypotheses described above could generalize to a diverse population of people. Participants were therefore selected from subject pools typically used in behavioral research. We reasoned that if there was strong support for a hypothesis in a broad population, then the hypothesis may be robust enough to investigate further in the population of interest (residents of high crime areas). Our primary participant pool was drawn from US users of Amazon.com's Mechanical Turk (MTurk) platform. For demographic information on MTurk participants, see: http://www.ipeirotis.com/wp-content/uploads/2017/12/wsdmf074-difallahA.pdf

participants would feel less comfortable in the less familiar scenes. But, we did not find any significant differences.[3] However, due to the lack of early experimental support for these hypotheses, we did not believe that they were robust enough to serve as the basis for a field intervention.

However, across a series of experiments, we found stronger and more consistent support for the hypothesis about officer anonymity. Recall that this hypothesis suggests that when people have more information about a police officer, they will believe that police officers has more information about them as well. In these online experiments, we tested a more general version of this hypothesis—namely that when people have more information about someone else (not just police officers), they will believe that person has more information about them as well. We will summarize three experiments that tested this hypothesis. In one experiment, 552 participants were told that they would be interacting with a stranger online. They were first asked to do an "icebreaker" in which they wrote down four things that were true about themselves and one lie about themselves. Participants were then randomly assigned to one of three conditions that varied the amount of information they believed were about another study participant: No information, seeing one truth about this other person, or seeing four truths. Note that there was no other participant. Instead, all interactions were simulated to make people believe that they were interacting with another person. Participants were then asked how likely it was that the other person could detect the participant's own lie. Participants who were given information about the stranger reported a higher probability that the stranger would detect their lie (approximately 37.5% versus 30.9% in the "no information" condition). That is, when participants were more familiar with the

[3] Note that because these experiments were far removed from the context in which crimes may happen (and with participants who may have a lower propensity to commit crimes), they cannot rule out the hypotheses above.

stranger, they thought that the stranger would better detect their lies (i.e., that the stranger had more knowledge about the participant).

We next tested whether this increased familiarity (and reduced anonymity) could *deter* dishonest acts. If increased familiarity can deter dishonest acts in a laboratory experiment, then it is possible that familiarity with police officers might deter crime in the real world. In another experiment, 588 participants from MTurk were told that they would be playing a game online in which they would flip a virtual coin 50 times. Each time they flipped a "head," they would earn some extra money. Participants were asked to enter the results of their coin flips into a text box. They were told that they were paired with another person—a judge who would determine whether they thought the participant was lying about how many heads they flipped. In one condition, participants were given no information about the judge. In another condition, they were given facts about the judge, such as what their favorite memories were or what a perfect day would feel like. Note that again there was no actual judge (interactions were simulated to make people believe that there was indeed another person who was serving as the judge). First, we found evidence that participants were somewhat dishonest—in both conditions participants reported flipping heads significantly more than 50% of the time. But this tendency was greater when participants did not have any information about the judge. Familiarity with the judge reduced the amount of dishonesty among participants.

It is natural to wonder about the ecological validity of the stylized experiments above. We therefore tested whether this hypothesis is relevant to policing. In partnership with community organizations, we recruited a convenience sample of 200 community members (from Brooklyn, Queens, the Bronx, and Manhattan) and 100 probationers (from Queens), and randomly assigned them to one of four conditions. In the "no information" condition, participants were simply told to imagine that there was a new officer (Officer Johnson) in their neighborhood. Other participants

were assigned to one of three information conditions: they read a letter from Officer Johnson explaining that he would be working in their community; they read a version of the letter with additional information about the officer (he liked pizza, he was a fan of the Knicks); they saw the officer's Facebook profile with the same information as the previous condition. Across a series of self-report scales, when participants knew more about the officer, they felt the officer knew more about them, were more likely to provide the officer with information about a crime, felt that the officer had more in common with their community and themselves, and trusted the officer more.

These results build on a series of findings in the psychology literature that also lend theoretical support to our hypothesis. For example, people often experience an illusion of transparency, or the sense that others can easily read one's internal states (Gilovich, Savitsky, & Medvec, 1998; Gilovich & Savitsky, 1999; Van Boven, Gilovich, & Medvec, 2003; Vorauer & Claude, 1998). For instance, people believe that their lies are easily detectable and emotional states are easily read (Gilovich et al., 1998). Second, people often anchor on their own perspective when making inferences about others (Epley, Keysar, Van Boven, & Gilovich, 2004; Nickerson, 1999; Savitsky, Keysar, Epley, Carter, & Swanson, 2011). If we anchor on how much we know about someone else, we might assume they know a similar amount about us. Third, when we have information about someone else, that reduces the perceived social distance from them (Small & Loewenstein, 2003, 2005), making it seem more plausible that they would also have more information about us.

Given these results, and the relevant psychological literature, we believed there was enough support for the hypothesis for it to inform the design and evaluation of a novel policing intervention. We tested this intervention in Phase 3.

PHASE 3: AN INTERVENTION TO REDUCE OFFICER ANONYMITY

This intervention was conducted with Neighborhood Coordination Officers (NCOs) who work the same shifts each week in defined geographic areas of a neighborhood with the goal of increasing NCOs familiarity with local residents and local problems. The NCOs are "off-radio" for a majority of their shifts and during this time they are tasked with engaging with neighborhood residents, identifying problems and working towards solutions. NCOs also work in and around public housing developments in New York City. We conducted this intervention with NCOs who were assigned to NYCHA developments prior to the intervention.

Based on surveys with the community and conversations with NCOs, we developed a 22 question survey that asked NCOs about small details from their lives (ranging from their favorite food to why they became an officer). Each NCO selected 3-5 questions they felt comfortable answering. Their responses were used to develop individualized outreach letters and cards.

To test the effects of the intervention, we randomly selected a subset of eligible NYCHA housing developments to receive the intervention. NYCHA developments were considered eligible for this intervention if their associated NCOs agreed to participate. While NCOs at 69 developments agreed to participate, these developments comprised 56 NCO/NYCHA development pairings ("development groups"), as some developments were grouped together due to close proximity and association with the same NCOs. Half of these 56 development groups were then randomly assigned to receive the intervention (treatment), while the other half were not (control).

This intervention involved two types of outreach to residents in treatment developments:

1. A mailer containing a letter from their local NCO (mailed to every apartment in the treatment NYCHA developments) with a brief introduction from the officer, along with a few innocuous

pieces of information about themselves and their contact information with an invitation to reach out if they need. The mailer also contained an "outreach card," designed as a physical reminder for residents, communicating the same personalized information about the NCO as the letter and contact information. The outreach card was sent with the introductory letter in three different waves of mailers between November 16, 2017 and January 6, 2018.

2. Participating NCOs conducted in-person outreach in treatment NYCHA developments, where they introduced themselves and distributed "outreach cards" to residents entering and leaving development buildings.

Following the intervention, we surveyed 1,960 individuals (18 years and older) across treatment and control developments. The surveys asked about police responsiveness in their area, officer familiarity with their area, residents' own levels of trust in the police, and other sentiments toward officers working in their area. The primary question of interest (based on our lab experiments) was whether the intervention would affect NYCHA residents' perception of police officers' awareness of them. Indeed, we found that residents in treatment developments thought it was more likely that officers in their neighborhood would find out if they did something illegal (we did not find a statistically significant effect on whether residents felt that officers knew them). Specifically, the intervention resulted in a 0.13 standard deviation increase in the average NYCHA resident's belief that an officer would find out if they committed a crime. This analysis provides support for our findings from the first phase of this project: that when people have more information about others, they believe that others have more information about them. This survey finding is important

because it suggests that light-touch interventions may be able to reduce officer anonymity, potentially acting to deter dishonest or criminal behavior.[4]

To evaluate the effects of the intervention on our primary crime outcomes of interest—criminal complaints and the arrest rate per reported crime—we used a combination of publicly- and non-publicly available NYPD data on crime committed on and near NYCHA developments[5] For on-campus crimes (i.e., those taking place on NYCHA property), results were not statistically significant, though suggestive that criminal complaints decreased for treatment relative to control developments, and that the arrest rate per reported crime increased for treatment developments.[6]

Given that several developments have few crimes occurring on campus, and that resident-related crimes can also happen just off campus, we also analyzed the effects of the intervention on near-campus crimes (i.e., crimes that take place within 250 feet of NYCHA development property).[7] The results for near-campus criminal complaints were similar in direction to the results for on-campus crimes, but slightly stronger and statistically significant. Specifically, the results suggest

[4] Separately, we found no evidence that the intervention had an effect on exploratory survey measures such as resident perceptions of police responsiveness, trust in police, or likelihood to report suspicious activity to the police.
[5] Non-publicly available data was provided by NYPD for all crimes filed by NYPD housing bureau officers that occurred on-campus for nine full months following the intervention, from Feb 2018 to Oct 2018. Publicly available NYPD data on all crimes and arrests on and near-campus, which include crimes filed by NYPD patrol bureau officers as well as crimes filed by NYPD housing bureau officers, has to date only been released for five full months following the intervention, from Feb 2018 to Jun 2018.
[6] These non-statistically significant results varied across specifications from a decrease of 0.3% to 9.6% for criminal complaints, and from an increase of 4.4% to 12.9% for arrest rate per reported crime.
[7] This distance was selected to include crimes occurring within one city block of each development, as 250 ft. is the average length of a NYC block. Note that, as we expected the intervention to *directly* affect resident perceptions of officers' awareness of them and only *indirectly* affect crime through this mechanism, we expected to have better statistical power to detect the effects on the survey mechanism than on crime. Defining the geography of crimes further complicates how to measure crime outcomes: setting them too small runs the risk of missing crimes committed by the target population, while geographies too large carry the risk of watering down our estimates by including crimes committed nearby by people not from the target population.

that the intervention led to a 10.6% decrease in near-campus criminal complaints, which translates to a 0.18 standard deviation decrease in criminal complaints.[8]

IMPLICATIONS

The results of this intervention have a striking implication. In many large cities, it is possible that individual officers remain relatively unknown to their communities. Police anonymity likely has several consequences. It may explain the growing reluctance among community members to come forward as witnesses. While not the only cause, officer anonymity likely contributes to the problem; few of us would share sensitive information with someone who is essentially a stranger. However, officer anonymity does not just affect whether people come forward with information. It may also make it more difficult to prevent crimes. If it is difficult for potential offenders to visualize how they could get caught or who the officers are who might apprehend them, then they may be more likely to believe that they are unlikely to be caught. Our intervention suggests that it might be possible to deter crime by familiarizing people with the officers in their community.

DISSEMINATION

We believe the results from this project could have important nationwide impacts on current understanding of neighborhood policing strategies and applications of behavioral science to policing. Throughout the duration of the project we sought to reach a broad audience of academic researchers, policymakers, and law enforcement practitioners through briefings, presentations, and regular updates. We include a complete list of dissemination activities in the Appendix.

[8] As an exploratory analysis, we broke down the results into particular types of crimes, and found suggestive evidence that the intervention decreased violation and outdoor crime complaints, though the effects were not statistically significant across either subgroup (across specifications and distances, the non-statistically significant results varied from a decrease of 1.8% to 23.8% for violations, and a decrease of 11.5% to 21.0% for outdoor crime complaints). Finally, we found no evidence that the intervention had an effect on other exploratory outcomes such as total arrests or the average time between criminal complaint occurrence and report.

CONCLUSION

We are grateful to the NIJ for its support of this work that we believe has generated insights into how to apply behavioral science to improve hot spot and neighborhood policing. We believe that this research will have a significant impact not only on policy, but also on the lives of individuals and communities in cities throughout the country.