

Predictive Policing and the Ethics of Preemption

DANIEL SUSSER

The American justice system, from police departments to the courts, is increasingly turning to information technology for help in identifying potential offenders; determining where, geographically, to allocate enforcement resources; assessing flight risk and the potential for recidivism among arrestees; and making other judgments about when, where, and how to manage crime. In particular, there is a focus on machine learning and other data analytics tools, which promise to accurately predict where crime will occur and who will perpetrate it. Activists and academics have begun to raise critical questions about the use of these tools in policing contexts. In this chapter, I review the emerging critical literature on predictive policing and contribute to it by raising ethical questions about the use of predictive analytics tools to identify potential offenders. Drawing from work on the ethics of profiling, I argue that the much-lauded move from reactive to preemptive policing can mean wrongfully generalizing about individuals, making harmful assumptions about them, instrumentalizing them, and failing to respect them as full ethical persons. I suggest that these problems stem both from the nature of predictive policing tools and from the sociotechnical contexts in which they are implemented. Which is to say, the set of ethical issues I describe arises not only from the fact that these tools are predictive, but also from the fact that they are situated in the hands of police. To mitigate these problems, I suggest we place predictive policing tools in the hands of those whose ultimate responsibility is to individuals (such as counselors and social workers), rather than in the hands of those, like the police, whose ultimate duty is to protect the public at large.

From Reactive to Preemptive Policing

Law enforcement has always utilized a mixture of reactive and proactive strategies, but several historical factors have conspired over the past few decades to give proactive approaches pride of place. First, rising crime rates in the 1960s and '70s led American police departments to question the largely reactive approaches of the middle part of the twentieth century, and shift toward what came to be known as “community policing” (Uchida 2005; Walsh 2005). Developments such as “broken windows” theory and problem-oriented policing in the 1980s and 1990s represented a change in “the definition of policing from one of crime control to one of community problem-solving and empowerment” (Uchida 2005, 36). That is, police departments began to focus their efforts on working with community members to transform the social and environmental conditions that encourage crime—in order to prevent it—rather than merely responding to criminal behavior.¹

The second historical factor was the rise of Compstat. Throughout the 1990s, William Bratton led the development of an approach to policing—first at the New York City Police Department, and then widely adopted elsewhere (Willis, Mastrofski, and Weisburd 2004)—which “builds upon the police organizational paradigms of the past and blends them with the strategic management fundamentals of the business sector” (Walsh 2005, 206). Central to that enterprise was the incorporation of huge amounts of data into law enforcement practice.² While the Compstat model retained community policing’s bottom-up orientation, allowing precinct commanders to determine local enforcement strategies, it introduced strict accountability from on high (Walsh and Vito 2004; Walsh 2005). To do so, the Compstat system required highly formalized data collection and processing to provide managers reliable, up-to-date information with which to conduct “relentless assessment” of their subordinates (Walsh and Vito 2004, 59; Willis, Mastrofski, and Weisburd 2007). But this turn to data-driven policing did more than facilitate new management processes; it enabled a change in focus from individual cases to larger crime patterns now visible in the data (Willis, Mastrofski, and Weisburd 2004), and a further entrenchment of the idea

that policing is about preventing crime rather than merely controlling it (Walsh 2005).

Finally, the terrorist attacks of September 11, 2001. Understood by many as being, at bottom, a failure of information sharing among intelligence and law enforcement agencies, 9/11 resulted in an even greater emphasis on data-driven approaches to crime control and the rise of “intelligence-led policing” (Ratcliffe 2008). As Sarah Brayne (2017) argues, after 9/11 local police came to be viewed as “actors on the front lines of the domestic war against terror,” leading local and federal law enforcement organizations to work together “to assess the viability of a more predictive approach to policing” (981). The emphasis on prediction and prevention over reaction and response was cemented in the process.

Alongside this historical and organizational background, in which police departments increasingly prioritized proactive policing strategies over reactive ones, one also finds the emergence of specific predictive tools in law enforcement practice. Although much of the emerging public discourse around predictive policing treats it as a fundamentally new phenomenon, police departments have long collected crime data, and they have long used it to make decisions about where limited resources ought to be devoted (Pearsall 2010; Perry et al. 2013).³ Until the advent of big data–driven analytics tools, this generally took the form of “crime mapping”—using information about where and what kinds of crimes have occurred in order to identify patterns useful for making forward-looking predictions about where criminal activity is likely to take place in the future. According to a report from the National Institute of Justice (NIJ) (the research arm of the US Department of Justice), this kind of “crime analysis” has been in use since at least the turn of the twentieth century, and perhaps longer (Harries 1999, 1), and was central to early iterations of Compstat (Willis, Mastrofski, and Weisburd 2004).⁴ Many will recognize the pre-digital form of this strategy in the physical maps on police department walls, with push-pins or other markers designating recent police activity, from films and television police procedurals.

As the NIJ report argues, crime mapping techniques developed piecemeal until the 1990s, when two concurrent trends marked a shift in their evolution: first, the incorporation of computers (especially geographic information systems, or GIS) into mapping and visualization practices, and second, the professionalization of crime analysis, includ-

ing the introduction of insights from academic criminology, geography, and cartography (Harries 1999, 3–6). Computerizing crime mapping offered several obvious advantages—it made it possible to analyze more information, to keep better track of trends (since old maps could be easily retained for comparison as new information was added, unlike what was possible with physical maps and push-pins), and as the cost of computers declined, more departments gained access to the tools and techniques of digital crime analysis (Harries 1999, 3–6). The professionalization of crime mapping, or what is sometimes called “crime forecasting” (Perry et al. 2013)—especially the incorporation of academic research into its ongoing development—meant that it could be driven by more than tacit knowledge (e.g., the beat cop’s familiarity with his or her route) and intuition, as social scientific theories designed to explain when, where, and why crime occurs were used to refine the analyses (Eck et al. 2005).⁵

Today, advanced crime mapping tools have been packaged into off-the-shelf products available for purchase by law enforcement agencies, without the need for in-house experts. PredPol and HunchLab are, perhaps, the most widely adopted and frequently discussed. Developed by anthropologist Jeffrey Brantingham and computer scientist George Mohler, PredPol analyzes data about when, where, and what kinds of crimes have occurred (focusing mainly on non-violent crime), in order to predict where hotspots are likely to emerge. As of early 2019, the software had been purchased and deployed in “dozens” of US municipalities (Haskins and Koebler 2019). HunchLab works in a similar way, though it implements an allegedly more advanced method of hotspot prediction known as “Risk Terrain Modeling,” which incorporates information about environmental factors, such as the presence of pawnshops and public transportation (which are correlated, for example, with higher burglary rates) into the predictions (Caplan and Kennedy 2011; Cham-mah 2016).

Crime mapping is fundamentally spatiotemporally oriented or place-based, the aim being to use information about historical crime patterns to identify “hot spots”—places and periods of time prone to above-average criminal activity (Eck et al. 2005; Perry et al. 2013). In addition to trying to predict where and when crime will occur, however, law enforcement agencies have also long attempted to understand *who* is likely

to be involved in criminal activity—either as a perpetrator or a victim. According to Andrew Ferguson (2016), such “person-based”⁶ predictive policing strategies developed independently of the spatiotemporally oriented varieties described above, emerging from public health research focused on crime. “For decades,” he writes, “sociologists identified the reality that a small subset of individuals in any community committed the vast majority of crimes. Police recognized that targeting those individuals could result in a disproportionate reduction of crime rates” (Ferguson 2016, 713).

This idea—that specific people, and groups of people, not just specific places and times—are statistically more likely to be implicated in criminal activity, led to the introduction of social network analysis into crime prediction. The premise of social network analysis is that sociological phenomena (such as crime) can be explained and predicted in part by the types and structures of social relationships that connect individuals (Borgatti et al. 2009). In the context of criminology, the relevant relationships or networks of interest have largely been gangs, drug dealing operations, and the mafia, all understood to be nexuses of crime. By analyzing the social networks of known offenders, police can identify others statistically likely to be involved in future criminal activity (Ferguson 2016; Tayebi and Glässer 2016).

As with crime mapping, the rise of cheap, powerful data-driven technologies meant new possibilities for social network analysis, as well as new insights into the correlations between other empirically observable phenomena and the potential for involvement in criminal activity. Contemporary person-based modeling “creates risk profiles for individuals in the criminal justice system on the basis of age, criminal record, employment history, and social affiliations” (Shapiro 2017). These profiles are then used to identify potential offenders and score their risk of involvement in future crimes. The Chicago Police Department, for example, maintains a “Strategic Subjects List” (or “heat list”) which contains information on thousands of citizens (Rosenblat, Kneese, and Boyd 2014; Dumke and Main 2017). Although the vast majority of people on the list have some kind of criminal record, not all do, and many (13 percent) have never been charged with a violent crime (Dumke and Main 2017). Person-based predictive policing technologies like Chicago’s heat list are, at the moment, less prevalent in the field than place-based tools

(Degeling and Berendt 2018), but their use is growing, having already been adopted by police departments in New Orleans, Los Angeles, Kansas City, and Rochester, New York (Ferguson 2017b).

These tools are changing how police perceive and interact with the public in important ways. First, as Sarah Brayne (2017) argues, data-driven policing means the police are able to effectively surveil an unprecedentedly large number of people. Network analysis, as we just saw, leads to people who have never had any contact with the police appearing in databases of potential offenders. And datasets from across public sector institutions are being integrated into police surveillance systems, meaning any contact with the state can potentially bring individuals to law enforcement's attention. Police officers no longer have to "actively search" for persons of interest, either. Instead, "passive alerts" notify officers whenever individuals matching certain profiles are found in the system. As Elizabeth Joh writes, "new technologies have altered surveillance discretion by lowering its costs and increasing the capabilities of the police to identify suspicious persons" (2016, 15).

In addition to preemptively identifying people for police scrutiny, predictive policing tools increasingly mediate encounters between individuals and police through risk assessments and threat scores. Intrado's "Beware" system, for example, claims to "help first responders understand the nature of the environment they may encounter during the window of a 9-1-1 event."⁷ Using information collected from public crime data, social media, and commercial data brokers, the system issues color-coded assessments—green, yellow, red—designed to indicate how dangerous an individual is to law enforcement (Robinson and Koepke 2016). As Ferguson argues, such assessments obviously

distort the day-to-day police decisions about use of force and reasonable suspicion. After all, once police have information that a person has a high threat score, this knowledge will color criminal suspicion and increase perceived danger, resulting in more frequent and more aggressive interactions with people the algorithm deems "high risk." (Ferguson 2017c)

Taken together, the historical trajectory of American law enforcement away from reactive and toward proactive approaches, and the introduction of increasingly data-driven, predictive tools into the ev-

eryday work of police, represent what Brayne calls “a migration of law enforcement operations toward intelligence activities” (2017, 986). “[L]aw enforcement,” she writes, “typically becomes involved once a criminal incident has occurred. . . . Intelligence, by contrast, is fundamentally predictive. Intelligence activities involve gathering data; identifying suspicious patterns, locations, activities, and individuals; and preemptively intervening based on the intelligence acquired” (Brayne 2017, 986). Obviously, such intelligence-driven, preemptive policing creates different risks and, when it goes wrong, threatens different harms than traditional, reactive law enforcement (Maguire 2000). And so this migration toward predictive, preemptive policing brings with it different ethical challenges.

Critical Responses to Predictive Policing

Activists and academics have raised a number of concerns about data-driven predictive policing techniques. First, many watching the emergence of predictive policing technologies and their integration into everyday policing practice worry, above all, about the lack of transparency around how these technologies work (what data they collect, the statistical methods used to analyze that data, etc.), how they are being deployed, and the concomitant absence of accountability (Dumke and Main 2017; Robinson and Koepke 2016). As Ferguson argues, this is in part simply the latest iteration of an old problem, given that “police accountability has long been a fraught issue” (2017a, 1168). But predictive technologies introduce new problematic dynamics: the tools are extremely complex, requiring special expertise to audit them; they are generally powered by proprietary algorithms, owned by private firms reluctant (and, so far, uncompelled) to reveal their inner workings; and they are still nascent technologies undergoing constant innovation and change (Ferguson 2017a). Such problems notwithstanding, calls for increased transparency around predictive policing technologies are growing (Robinson and Koepke 2016; Schmidt 2018; Zarsky 2013), mirroring concerns about algorithmic opacity in other contexts, such as credit scoring and search (Pasquale 2015).

Second, critics have raised questions about accuracy and bias. Worries about bias in predictive policing technologies are a special case of

concerns about “algorithmic bias” more generally, an issue stemming from realizations—beginning in the mid-1990s—that computers can encode and perpetuate discriminatory effects on their users and others impacted by their operations (Friedman and Nissenbaum 1996). More recently, it has become apparent that the class of computational techniques known as machine learning—which drives a majority of predictive tools—is especially vulnerable to this problem.⁸ Because machine learning tools “learn” to make predictions from patterns latent in the data they are trained on, their predictions carry forward any biases or other effects of discrimination reflected in that data. Given the history of discrimination against Black and minority populations in the United States, predictive policing algorithms trained on historical data are likely to perpetuate and reinforce racial biases (Ferguson 2017a; Kutnowski 2017; Selbst 2017; Shapiro 2017). And indeed, that is what we have already begun to see. An algorithm for predicting potential recidivism, commonly used in courts around the US to make bail, sentencing, and parole decisions, is more likely to erroneously flag Black offenders as dangerous than White offenders (Angwin et al. 2016), and place-based predictive algorithms used by police departments to determine where and when to deploy officers results in the disproportionate over-policing of majority Black neighborhoods (Lum and Isaac 2016).

A third—related—strand of critique looks not at bias in training data, but rather at the feedback loops and multiplier effects that can result from putting predictive algorithms into practice. When predictive policing tools are deployed in the field, they continue to collect data, which is then used to update and refine their predictive models. As a result, the models become trained over time on a disproportionate amount of information about the people and places they directed law enforcement to in the first place, producing feedback loops (Ensign et al. 2018). People and places historically discriminated against thus become targets of police attention (leading to more arrests, etc.), which leads to data that appears to confirm the original predictions. As Kristian Lum and William Isaac write, “newly observed criminal acts that police document as a result of these targeted patrols then feed into the predictive policing algorithm on subsequent days, generating increasingly biased predictions. This creates a feedback loop where the model becomes increasingly confident that the locations most likely to experience further

criminal activity are exactly the locations they had previously believed to be high in crime: selection bias meets confirmation bias” (2016, 16; see also Ferguson 2017a).⁹

Finally, some have asked not about the predictions of predictive policing, nor about the enforcement dynamics the predictions produce, but rather about how predictive technologies change the way people think about and understand the role of law enforcement generally. Bernard Harcourt articulated an early version of this worry in *Against Prediction* (2007). Tracing the history of what he calls “actuarial approaches” to policing (meaning predictive approaches grounded in statistical methods), he argues that beyond concerns about the accuracy of statistical models or the utility of such models for actually preventing crimes (concerns he also raises), a more fundamental problem with incorporating predictive tools into law enforcement practice is the effect it has on our conception of the law enforcement enterprise. Just as holding a hammer makes everything look like a nail, reliance on statistical methods for structuring the work of law enforcement reshapes our understanding of that work in its actuarial image. “We have come to associate the prediction of future criminality with just punishment,” Harcourt writes. “This seems intuitively obvious, even necessary. . . . But, the fact is, we have chosen this conception of just punishment. . . . We chose it as against a rehabilitative model and as against a more strictly retributivist model” (2007, 31–32). As the work of law enforcement has come to be understood fundamentally in terms of predicting and mitigating risk, the rationale for subjecting people to law’s coercive force has likewise reoriented toward people’s predicted riskiness.¹⁰

In what follows, I continue in the vein of this last strand of critique, raising questions about the effects of predictive technologies on the treatment of individuals by the state. Specifically, I attempt to shed light on some ethical ramifications of person-based predictive policing—law enforcement officials singling out individuals for investigative scrutiny on the basis of algorithmically generated predictions. I leave questions about the ethics of place-based predictive policing to the side.

Problems with Profiling

As we’ve seen, person-based predictive policing is, in effect, a data-driven form of profiling. Unlike the beat cop evaluating someone’s suspicious

behavior in the moment, predictive policing tools evaluate individuals on the basis of whether they are a certain type—the *kind of person* who is likely to commit a crime.¹¹ Using information about who has committed past crimes¹² (their age, employment history, social networks, and so on), law enforcement generates “risk profiles”—assessments of how similar individuals are to past perpetrators, and thus how likely they are to engage in future criminal activity (Shapiro 2017). While the technology driving this kind of profiling is new, the underlying idea is not. Law enforcement has long engaged in low-tech forms of profiling, and scholars and advocates have raised concerns about the ethical ramifications of such strategies. To understand the ethics of person-based predictive policing, it therefore helps to consider arguments about the ethics of profiling more generally.

Profiling and Generalization

First, it is perhaps important to distinguish the sort of profiling at issue here from what has traditionally been referred to as “criminal profiling”—i.e., the use of crime scene evidence to infer the perpetrator’s personality in order to guide an investigation (Gregory 2005; Turvey 2012). Such profiling works backward from an actual crime to an actual criminal (rather than a potential crime and potential criminal). Though data analytics tools are today frequently used to aid criminal profilers, and such tools are sometimes described using the umbrella term “predictive policing,” they raise substantially different ethical questions than those raised by the predictive tools that identify individuals likely to commit future crimes (Perry et al. 2013).¹³ My argument in what follows does not bear on criminal profiling, at least directly.

A more familiar and more relevant strategy is *racial profiling*—selecting people for differential treatment on the basis of their race. Racial profiling, especially in the context of law enforcement, has long been a subject of controversy. Discussions about whether it is morally acceptable for law enforcement to engage in racial profiling are relevant to debates about the ethics of predictive policing for at least two reasons. First, both racial profiling and person-based predictive or algorithmic profiling involve *preemptive* interventions by law enforcement. For reasons discussed above, police departments have become increasingly

focused on preventing crime, rather than merely responding to it. The purpose of both racial profiling and predictive profiling is to determine in advance of any outwardly suspicious behavior whether or not particular individuals are especially likely to commit a crime, and to intervene before they do. Second, in both cases the *method* of making such determinations is statistical inference. Which is to say, in the absence of outwardly suspicious behavior, judgments about individual risk are made by examining the statistical correlations between an individual's group memberships and criminality more generally.

Examples of racial profiling by law enforcement abound. Customs agents have disproportionately subjected African American women to invasive treatment at US airports (Schauer 2003). In cases of so-called “pretext stops,” Black drivers have been disproportionately stopped and searched on US roads (Hosein 2018). Arab and Muslim people in the US—especially after 9/11—have been targeted with suspicion simply because of their race or religion, and they have been singled out for questioning and other forms of differential treatment by law enforcement and intelligence agencies (Hosein 2018). Such practices are so prevalent that in 2009 the American Civil Liberties Union reported that “the practice of racial profiling by members of law enforcement at the federal, state, and local levels remains a widespread and pervasive problem throughout the United States, impacting the lives of millions of people in African American, Asian, Latino, South Asian, and Arab communities” (American Civil Liberties Union and The Rights Working Group 2009).

Like the predictive or algorithmic profiling described above, racial profiling is a form of “statistical discrimination” (Lippert-Rasmussen 2014; Schauer 2018). Which is to say, it involves sorting or classifying people on the basis of observable traits assumed to be statistically correlated with some target characteristic. In the case of intelligence and law enforcement agencies racially profiling Arab Americans, for example, it is assumed that race or ethnicity is meaningfully statistically correlated with the propensity for terrorist acts. Because what law enforcement is interested in—the disposition to terrorism—is not easily observable, race is used as a proxy. Person-based predictive policing tools work the same way. What law enforcement wants to know—whether someone is especially likely to commit a crime—cannot be directly observed, so

proxy attributes are used instead. Age, employment status, past criminal charges, and other characteristics statistically correlated with criminality are used as the basis for discriminating between those who pose the least and those who pose the greatest risks for future criminal behavior.

“Discrimination” has a pejorative connotation, and racial profiling has often been criticized for being discriminatory. As Frederick Schauer argues, however, discrimination cannot itself be the problem, because acting on the basis of statistical inferences of the sort just described is unavoidable. At a basic level, to discriminate simply means to make value-laden distinctions. We discriminate when we choose to eat tasty rather than bland foods, when we go for walks on clear rather than rainy days, and when we hire charismatic rather than lackluster job candidates. And making these kinds of distinctions almost always requires statistical inferences. Having previously found a certain brand of food bland, we infer that it is very likely the same brand will be bland this time. Because we found that previous job candidates who lacked charisma during their interviews were ultimately unmotivated in their positions, we assume that this uncharismatic job candidate would be unmotivated in this position too. Inductive reasoning—*making generalizations*—is fundamental to human judgment (Schauer 2003).

That racial profiling seems to many intuitively morally unacceptable therefore cannot be explained simply by the fact that it involves statistical discrimination.¹⁴ Rather, for Schauer, racial profiling is ethically suspect because it involves *bad statistics*. In most cases of profiling on the basis of race (or, equally, gender, sexual orientation, or other social categories), the alleged correlations between such categories and most target characteristics simply turn out to be spurious (Schauer 2003).¹⁵ Even though some people think women, as a group, lack business acumen, or gay men lack physical courage—to take Schauer’s examples—there is no justification for those beliefs (Schauer 2003). Such generalizations (or “stereotypes”) about social groups are just wrong. They are *prejudices*, not sound inductive inferences. Judging individuals on the basis of false generalizations about the social groups of which they are members is wrong, because it is highly likely that the judgments will be mistaken. As Schauer writes, “much of the history of unfortunate discrimination is a history of the erroneous belief in statistical relationships that turn out to have no basis in fact” (Schauer 2018, 47).

Picking Proxies

But false generalizations are an easy case. What should we make of statistically well-grounded profiling? Is it acceptable to judge individuals on the basis of *nonspurious* correlations between group membership and target characteristics? If, in theory, race really was correlated with rates of criminality, would racial profiling be justified? Is it reasonable to assume that someone poses an especially high risk for future criminal behavior just because they share certain traits in common with other criminals? Some argue that the answer is categorically no. Even in cases where there is a high degree of certainty that a generalization about someone will turn out to be true, it is wrong to make consequential decisions about them on that basis, because the chance that the judgment is mistaken—however small—denies them the right to individualized treatment (see Schauer 2003, 20–22; Lippert-Rasmussen 2011, 2014, 275). The moral intuition behind this categorical rejection of statistical discrimination—that individuals ought not to be judged using generalizations about groups to which they belong, but rather only on the basis of their own particular characteristics—is powerful and widely shared.

As Schauer (2003) argues, the idea of individualized or particularized judgment is implausible when taken to its extreme, because even judgments that seem on the face of it highly individualized are usually, at bottom, grounded in statistical inferences.¹⁶ However, critics of racial profiling need not argue that all profiling is unacceptable to demonstrate that racial profiling is. They only need to show that it is unacceptable to use race, specifically, as a proxy variable when making certain kinds of judgments, such as whom to single out for police attention. Kasper Lippert-Rasmussen (2014) adopts this strategy, arguing that using race as a proxy for criminality is wrong—even if it is sometimes statistically sound—because the history of discrimination against racial minorities is likely an important cause of disparate crime rates between races, and using that statistical disparity as a justification for burdening minority groups offends norms of fairness. Adam Hosein (2018) argues it is wrong because using race as a proxy for criminality contributes to a sense of inferior political status among targeted racial minorities.

Returning to the question of predictive policing, Schauer's argument (that race is usually a poor proxy for other traits) and Lippert-Rasmussen's and Hosein's arguments (that even when race is a sound proxy it is probably wrong to use it) demonstrate that while statistical inference may be unavoidable, there are important ethical questions to consider about *how such inferences are made*. At the very least, information about race should be excluded from the datasets predictive policing algorithms analyze when generating risk scores. Moreover, since such arguments likely apply equally to other important "socially salient groups" or "protected attributes" besides race—such as gender, sexual orientation, and ability—we can reasonably conclude that data about membership in those groups ought to be excluded too (Lippert-Rasmussen 2014). And to some extent, this is what we find. Although, as we've seen, there is insufficient transparency around most existing predictive policing algorithms, some law enforcement agencies—such as the Chicago Police Department—insist that they do not use data about race to generate their scores (Dumke and Main 2017; Stroud 2014).

There is reason to worry, however, that this solution—excluding certain kinds of information from predictive policing algorithms—while laudable, is insufficient. Namely, it is a feature of machine learning algorithms (which are, again, the computational technique underlying most predictive policing systems) that they will find proxies for meaningful variables regardless of whether or not those variables are explicitly introduced into the system. As Solon Barocas and Andrew Selbst write: "Even in situations where data miners are extremely careful, they can still effect discriminatory results with models that, quite unintentionally, pick out proxy variables for protected classes" (2016, 675). Indeed, the veneer of mathematical objectivity offered by algorithmic assessments can in fact serve to mask the ways such assessments further entrench racial bias in policing. "Mathematized police practices," writes Brayne, "serve to place individuals already under suspicion under new and deeper forms of surveillance, *while appearing to be objective*, or, in the words of one captain, 'just math'" (2017, 997, emphasis in original). Thus, predictive profiling doesn't merely raise similar ethical questions to those raised by racial profiling, it threatens to reintroduce racial profiling in a new guise.¹⁷

Past as Prologue

Social group membership is not the only kind of proxy variable deserving of critical scrutiny. The language of “predictive policing” evokes a sense of futurity, an orientation toward what has yet to come. But, of course, predictions must be generated out of something, and that something is data about the past.¹⁸ The people on the Chicago Police Department’s heat list are there not because of anything they are doing in this moment; the list is updated frequently, but not in real time. They are on the list because they have been charged with crimes in the past, or they have been the victims of crime, or because they have been associated with known offenders, or because they are a certain age (Dumke and Main 2017). They may not have done anything remotely indicative of potential criminality in some time—for those on the list simply because they were the victims of crime, they may *never* have acted suspiciously—and yet the Strategic Subjects List indicates to law enforcement that they are a risk, a potential threat, and that they ought to be treated as such.

As we’ve seen, this is part of a larger trend away from reactive policing postures, toward preventive or preemptive ones. Prior to the emergence of intelligence-led policing, individuals were generally singled out for police attention because they were engaged—in the moment—in suspicious acts (Brayne 2017). Racial profiling is an exception to that rule, and person-based predictive policing technologies encourage another. Chicago’s Strategic Subjects List and other systems like it enable police to identify purportedly “high-risk” individuals by comparing them with known criminals who have similar biographies—similar criminal records, similar histories of criminal victimization, past associations with similar people. Which is to say, they allow the police to generalize about individuals using not race but history as a proxy for potential criminality. Once high-risk individuals are identified, they are subjected to differential treatment. In the Chicago case, this means visits from teams of police officers warning that they are under extra scrutiny (Gorner 2013; Stroud 2014). As we saw above, it can also mean police officers in the field approaching purportedly high-risk individuals with greater vigilance (Ferguson 2017c).

This kind of predictive strategy—using information about people’s past associations, past behaviors, past run-ins with the law to make pre-

dictions about their risk of future criminality—freezes people, Medusa-like, in their pasts, suspending them, condemning them to permanent stigmatization. Once identified by the system as a potential threat (a “strategic subject”), a person’s history is set in stone and they are marked by the state as inalterably suspicious. As Bonnie Sheehey argues, the predictive system

codifies and stabilizes the past, turning it both into something that is bound to repeat in the future and into something that can be securely acted on in the present. This codification of the past in the form of data functions to close the past off from possibilities of what could have been. No longer open or negotiable, the past gets preformed and packaged in the shape of data as something already given. (Sheehey 2018, 8)

Using a person’s history to mark them as suspicious in the present renders them the “type” of person who would, of course, have the kind of historical record they do. (As we saw in the case of place-based predictive policing: “selection bias meets confirmation bias.”) In doing so, their very identity becomes cause for suspicion—they become *suspicious people*. They need not engage in any overtly suspicious behavior to attract attention from law enforcement, and, indeed, there is nothing they can do (or abstain from doing) to avoid such attention.¹⁹

These considerations suggest additional ethical concerns about police profiling, beyond those raised in the previous section. First, by marking individuals as suspicious people, regardless of any overtly suspicious behavior, predictive technologies rob them of the presumption of innocence. Legally, in the US, that presumption is only formally required in court once someone has been charged with a crime.²⁰ Morally, however, the right to be presumed innocent is broader than that. As Ian Kerr and Jessica Earle (2013) argue,

If the legal universe has a prime directive, it is probably the shared understanding that everyone is presumed innocent until proven guilty. . . . [T]he presumption of innocence and related private sector due process values can be seen as wider moral claims that overlap and interrelate with core privacy values. (Kerr and Earle 2013)

The kind of “preemptive predictions” (to use Kerr and Earle’s term) predictive policing technologies generate—which is to say, predictions that serve to anticipate and constrain people’s behavior—undermine bedrock moral and legal values designed to “provide fair and equal treatment to all by setting boundaries around the kinds of assumptions that can and cannot be made about people” (Kerr and Earle 2013). Once someone is flagged by a person-based predictive policing algorithm, the police are authorized to make assumptions about them that they are forbidden from making about others.

At a deeper level, this treatment demonstrates a failure to recognize targeted individuals as full moral subjects. A hallmark of moral personhood is the capacity for agency—the ability, precisely, to thwart predictions, to surprise. Presuming to know how individuals will behave in the future, and preemptively, punitively intervening on the basis of that assumption suggests one of two things: either that law enforcement is sure they are right about how targets will act in the future, or they think the benefits of guessing correctly outweigh the costs of getting it wrong. If the former is true, then the deployment of predictive policing technologies treats individuals as rigid, mechanical automata, rather than as dynamic, agential beings—that is, as moral subjects. If the latter is true, then the deployment of predictive policing technologies acknowledges that its targets are moral subjects but then proceeds to *instrumentalize* them, treating their interests as subservient to the interests of others. Whether the truth is closer to the former or the latter, its ethical prospects are poor.

Of course, some will object that given the existence of predictive profiling tools, and the significant threat of crime, it would be negligent for the police *not* to use them. As I suggest in closing, however, there is, in fact, a third option: the power of predictive profiling technologies could be harnessed by others besides the police.

Conclusion: Just Preemption

Philosophers and science and technology studies scholars have long argued that in order to evaluate the social and ethical implications of new and emerging technologies, it is not enough to investigate the

technologies themselves; one must explore the broader sociotechnical contexts in which technologies are developed and used (see, e.g., Latour 2005; Pinch and Bijker 1984; Winner 1980). The question is not simply, “Is this technology good or bad? Does it promise to make things better or worse?” Rather, the question must be, “What does it mean for this technology to be developed and used in this way, in this context, by these particular social actors?”

I want to conclude by suggesting that the problems I have pointed to throughout this essay arise not simply from the predictive dimension of predictive policing technologies, but also from the fact that these technologies are situated in the hands of police. The duty of police officers is not only to individuals, but also to the public at large, and as a result, the work of policing is not fundamentally oriented to the *interests* of individuals, but rather to the public interest. Scholars of policing have shown that police work is far more complex and multifaceted than merely “enforcing the law,” and many police officers spend much of their time aiding individual citizens (see Bittner 1970; Goldstein 1977). But the principal goal of policing, its basic function, is ensuring public safety.²¹ Advocates of community policing attempted to “rearrange priorities” among various police functions in the 1990s, de-prioritizing crime fighting and elevating individual assistance (Eck and Rosenbaum 1994). As we saw above, however, the trend over the last two decades has been to reestablish crime prevention as law enforcement’s primary mission.²²

It is in this context that the outputs of predictive policing algorithms—risk assessment scores, strategic subjects lists—are interpreted and acted upon. To be designated “high-risk” or a “strategic subject” *means* that one is a threat, a potential obstacle standing in the way of public safety. And it is this understanding that shapes the particular preemptive interventions made on the basis of such designations. Being identified as high risk or a strategic subject means facing heightened vigilance on the part of police (which can be dangerous to the person deemed a threat), or a visit from police officers warning that one is under heightened scrutiny (which can be embarrassing and demeaning, and a dangerous signal to others).²³ The worries discussed in this chapter, about individualized judgment and racial stigmatization and the presumption of innocence

and potential instrumentalization, all carry normative force because it is police officers—people duty bound to promote the public interest over and above the interests of individuals, and authorized to use coercive force in carrying that duty out—who are making decisions on the basis of these predictions.

That these worries are in part a function of context suggests that they might be diffused if predictive profiling technologies were otherwise situated.²⁴ It is easy to imagine such tools driving the work of social services agencies, for example, rather than law enforcement.²⁵ In such a context, the prediction that someone was likely to commit a crime would mean something very different. Since the work of those who provide social services—such as social workers, counselors, and job placement officers—is oriented primarily toward the needs and interests of individuals (and only indirectly toward the public interest), to them a person appearing on a strategic subjects list is not, in the first place, a threat, but rather someone in need of assistance.

If predictive profiling technologies were placed in the hands of social workers, or anyone else whose work is fundamentally oriented to each individual's needs and interests, consider how the ethical questions raised throughout this essay would change: Is it unfair to generalize about someone in order to determine whether or not they need extra care? Are we comfortable treating someone as an underserved person, irrespective of any outwardly observable demonstrations of need, and to intervene preemptively to provide them with services? Is it acceptable to preemptively offer a person extra services, even at the risk that our predictions about their needs are proven wrong? By shifting the context in which predictive technologies guide people's work, we shift the moral calculus. Rather than asking if and when it is right to make individuals absorb the costs of erring on the side of public safety, we ask if and when it is right to make the public absorb the costs of erring on the side of individual welfare.

The answer to this last question is not obviously “yes, always.” A world in which predictive profiling tools were given to social workers rather than the police would bring with it its own ethical tradeoffs and complexities. But that hypothetical world will likely strike many as more plausibly just than our own world, where such tools are controlled by the police.

NOTES

- 1 Community policing “seeks to turn patrol from a reactive to a proactive function” (Walsh 2005, 205).
- 2 Willis, Mastrofski, and Weisburd refer to Compstat as a “data-saturated environment” (2007, 148).
- 3 Perry et al. (2013, xiv) write: “We found a near one-to-one correspondence between conventional crime analysis and investigative methods and the more recent ‘predictive analytics’ methods that mathematically extend or automate the earlier methods.”
- 4 Jennifer Bachner (2014) points to examples of crime visualization from as far back as 1829 in France, and claims that crime analysis of this sort was in regular use in mid-nineteenth century England, predating the establishment of formal police forces in the United States.
- 5 A number of theoretical frameworks have been used to support place-based crime forecasting, from “repeat victimization theory” to “routine activity theory” to “broken windows theory,” each pointing to a different set of environmental factors correlated with criminal activity. See Eck et al. (2005).
- 6 This is sometimes referred to as “offender-based predictive policing.” I prefer to follow Robinson and Koepke (2016) in calling it “person-based,” as many of those identified by such tools turn out not to be offenders at all.
- 7 From Intrado’s website (www.intrado.com/beware), quoted in Joh (2016, 24).
- 8 An entire field has quickly emerged in response to the problem of bias in machine learning algorithms. Work presented at its flagship conference, “FAT*: ACM Conference on Fairness, Accuracy, and Transparency,” gives a sense of the scope of the problem and methods for remedying it. See <https://fatconference.org>.
- 9 Bernard Harcourt (2007) locates this dynamic in pre-machine learning predictive tools as well; he calls it the “Ratchet Effect.”
- 10 Sheehey (2018) raises related questions.
- 11 A different class of automated technologies, which Michael Rich (2016) calls “Automated Suspicion Algorithms,” do in fact examine behavior in real time, in addition to historical, demographic, and social network information.
- 12 Or more precisely, who has been *charged* with committing past crimes.
- 13 For a discussion of ethical issues in criminal profiling, see Boylan (2011) and Turvey (2012).
- 14 For a qualified defense of racial profiling, see Risse and Zeckhauser (2004).
- 15 However, Schauer defends the use of profiling in many cases where the statistical correlation between the proxy attribute and the target attribute holds.
- 16 Even judgments based on direct observations of individuals are usually made by comparing what one is observing, probabilistically, with previous, similar observations.
- 17 Hanni Fakhoury, a staff attorney at the Electronic Frontier Foundation, said about this: “It ends up being a self-fulfilling prophecy. The algorithm is telling you exactly what you programmed it to tell you. ‘Young black kids in the south side of Chicago

are more likely to commit crimes,' and the algorithm lets the police launder this belief. It's not racism, they can say. They are making the decision based on what the algorithm is, even though the algorithm is going to spit back what you put into it. And if the data is biased to begin with and based on human judgment, then the results the algorithm is going to spit out will reflect those biases" (Llenas 2014).

- 18 According to the *New York Times*, the Chicago Police Department's Strategic Subjects List "draws, the police say, on variables tied to a person's past behavior, particularly arrests and convictions, to predict who is most likely to become a 'party to violence'" (Davey 2017).
- 19 Interestingly, this dynamic arguably exhibits the logic of racialization. Beyond the fact that the machine learning algorithms driving these systems very likely identify and incorporate proxies for race in their calculations—i.e., despite best efforts, race is likely an *input*—race is also, in some sense, an *output*. These systems create classificatory hierarchies of relative risk and they subject members of the different classificatory groups to systematically different treatment. Those deemed unthreatening are left alone; those deemed risky are subjected to heightened surveillance and the promise of more severe punishment if found engaging in criminal behavior.
- 20 Predictive policing is not "pre-crime"—the law enforcement bureau from Philip K. Dick's "Minority Report," that incarcerates people for crimes they are predicted to, but have not yet, committed—though some argue that predictive policing moves us in that direction (see McCulloch and Wilson 2016).
- 21 Jerome Hall writes that "organized police forces have functioned everywhere and at all times to maintain order principally by preventing crimes and apprehending offenders" (1953, 139). Walsh and Vito claim that "Police departments are created to provide public safety for a defined governmental jurisdiction" (2004, 51).
- 22 Indeed, as Johnny Nhan writes, it is characteristic of contemporary "police culture" that "Most officers marginalize activities that are not oriented toward law enforcement as not 'real' police work and consider such peacekeeping activities as a form of 'social work'" (2013, 5). See also Sierra-Arévalo's contribution in this volume.
- 23 A *Chicago Tribune* article describes one person's reaction to being identified as a strategic subject: "Interviewed at his Austin home, McDaniel said he was offended at being singled out by West, commander of the Austin police district. All the attention made him nervous because his neighbors noticed, leading them, he feared, to wonder if he was a police snitch. Two officers waited outside on the porch while the commander and a criminal justice expert spoke to McDaniel in his home. 'Like I said, I have no (criminal) background, so what would even give you probable cause to watch me?' said McDaniel, a high school dropout. 'And if you're watching me, then you can obviously see I'm not doing anything'" (Gorner 2013).
- 24 Kutnowski (2017) makes a similar suggestion about place-based predictive policing tools, arguing that they should reside not with police but rather with policy makers and urban planners. But he offers a different rationale.

- 25 Indeed, in Chicago, the police often bring social workers with them when they visit strategic subjects, or they provide them with information about how to access social services. Because the overall context is law enforcement, however, the goal of treating the underlying causes of criminal behavior necessarily takes a back-seat role to the goal of deterring it. For example, as Ferguson points out (quoting Andrew Papachristos and David Kirk), when Chicago police and prosecutors visit gang members on the Strategic Subjects List, “the point of the message stresses the deterrent aspect of the program,” and informs the target that, if they are caught misbehaving, “you will be punished harder because you were warned” (Ferguson 2016, 718–19).

BIBLIOGRAPHY

- American Civil Liberties Union and The Rights Working Group. 2009. *The Persistence of Racial and Ethnic Profiling in the United States: A Follow-up Report to the U.N. Committee on the Elimination of Racial Discrimination*. www.aclu.org.
- Angwin, Julia, Jeff Larson, Surya Mattu, and Lauren Kirchner. 2016. “Machine Bias.” *ProPublica*, May 23, 2016. www.propublica.org.
- Bachner, Jennifer. 2014. “Predictive Policing: Preventing Crime with Data and Analytics.” IBM Center for the Business of Government.
- Barocas, Solon, and Andrew D. Selbst. 2016. “Big Data’s Disparate Impact.” *California Law Review* 104 (3): 671–732.
- Bittner, Egon. 1970. *The Functions of the Police in Modern Society: A Review of Background Factors, Current Practices, and Possible Role Models*. Chevy Chase, MD: National Institutes of Mental Health.
- Borgatti, Stephen P., Ajay Mehra, Daniel J. Brass, and Giuseppe Labianca. 2009. “Network Analysis in the Social Sciences.” *Science* 323 (5916): 892–95.
- Boylan, Michael. 2011. “Ethical Profiling.” *Journal of Ethics* 15 (1–2): 131–45.
- Brayne, Sarah. 2017. “Big Data Surveillance: The Case of Policing.” *American Sociological Review* 82 (5): 977–1008.
- Caplan, Joel M., and Leslie Kennedy W., eds. 2011. *Risk Terrain Modeling Compendium: For Crime Analysis*. Rutgers, NJ: Rutgers Center on Public Security.
- Chammah, Maurice. 2016. “Policing the Future.” *The Verge*, February 3, 2016. www.theverge.com.
- Davey, Monica. 2017. “Chicago Police Try to Predict Who May Shoot or Be Shot.” *New York Times*, December 21, 2017. www.nytimes.com.
- Degeling, Martin, and Bettina Berendt. 2018. “What Is Wrong about Robocops as Consultants? A Technology-Centric Critique of Predictive Policing.” *AI & SOCIETY* 33 (3): 347–56.
- Dumke, Mike, and Frank Main. 2017. “A Look Inside the Watch List Chicago Police Fought to Keep Secret.” *Chicago Sun-Times*, May 18, 2017. https://chicago.suntimes.com.
- Eck, John E., Spencer Chainey, James G. Cameron, Michael Leitner, and Ronald E. Wilson. 2005. *Mapping Crime: Understanding Hot Spots*. Washington, DC: National Institute of Justice.

- Eck, John E., and Dennis P. Rosenbaum. 1994. "The New Police Order: Effectiveness, Equity, and Efficiency in Community Policing." In *The Challenge of Community Policing: Testing the Promises*, edited by Dennis Rosenbaum, 3–24. Thousand Oaks, CA: SAGE Publications.
- Ensign, Danielle, Sorelle A. Friedler, Scott Neville, Carlos Scheidegger, and Suresh Venkatasubramanian. 2018. "Runaway Feedback Loops in Predictive Policing." *Proceedings of Machine Learning Research* 81:1–12.
- Ferguson, Andrew Guthrie. 2016. "Predictive Prosecution." *Wake Forest Law Review* 51 (3): 705–44.
- Ferguson, Andrew Guthrie. 2017a. "Policing Predictive Policing." *Washington University Law Review* 94 (5): 1109–89.
- Ferguson, Andrew Guthrie. 2017b. *The Rise of Big Data Policing: Surveillance, Race, and the Future of Law Enforcement*. New York: New York University Press.
- Ferguson, Andrew Guthrie. 2017c. "Police Are Using Algorithms to Tell Them If You're a Threat." *Time*, October 3, 2017. <http://time.com/>
- Friedman, Batya, and Helen Nissenbaum. 1996. "Bias in Computer Systems." *ACM Transactions on Information Systems* 14 (3): 330–47.
- Goldstein, Herman. 1977. *Policing a Free Society*. Cambridge, MA: Ballinger.
- Gorner, Jeremy. 2013. "Chicago Police Use 'Heat List' as Strategy to Prevent Violence." *Chicago Tribune*, August 21, 2013. www.chicagotribune.com.
- Gregory, Nathan. 2005. "Offender Profiling: A Review of the Literature." *British Journal of Forensic Practice* 7 (3): 29–34.
- Hall, Jerome. 1953. "Police and Law in a Democratic Society." *Indiana Law Journal* 28 (2): 133–77.
- Harcourt, Bernard E. 2007. *Against Prediction: Profiling, Policing, and Punishing in an Actuarial Age*. Chicago: University of Chicago Press.
- Harries, Keith. 1999. *Mapping Crime: Principle and Practice*. Washington, DC: National Institute of Justice.
- Haskins, Caroline, and Jason Koebler. 2019. "Dozens of Cities Have Secretly Experimented with Predictive Policing Software." *Motherboard*, February 6, 2019. <https://motherboard.vice.com/>
- Hosein, Adam Omar. 2018. "Racial Profiling and a Reasonable Sense of Inferior Political Status." *Journal of Political Philosophy* 26 (3): e1–e20.
- Joh, Elizabeth E. 2016. "The New Surveillance Discretion: Automated Suspicion, Big Data, and Policing." *Harvard Law & Policy Review* 10 (1): 15–42.
- Kerr, Ian, and Jessica Earle. 2013. "Prediction, Preemption, Presumption." *Stanford Law Review Online* 66 (September): n.p. www.stanfordlawreview.org/
- Kutnowski, Moish. 2017. "The Ethical Dangers and Merits of Predictive Policing." *Journal of Community Safety & Well-Being* 2 (1): 13–17.
- Latour, Bruno. 2005. *Reassembling the Social: An Introduction to Actor-Network-Theory*. Clarendon Lectures in Management Studies. New York: Oxford University Press.

- Lippert-Rasmussen, Kasper. 2011. "‘We Are All Different’: Statistical Discrimination and the Right to Be Treated as an Individual." *Journal of Ethics* 15 (1–2): 47–59.
- Lippert-Rasmussen, Kasper. 2014. *Born Free and Equal? A Philosophical Inquiry into the Nature of Discrimination*. New York: Oxford University Press.
- Llenas, Bryan. 2014. "Brave New World of ‘Predictive Policing’ Raises Specter of High-Tech Racial Profiling." *Fox News*, February 25, 2014. www.foxnews.com/
- Lum, Kristian, and William Isaac. 2016. "To Predict and Serve?" *Significance* 13 (5): 14–19.
- Maguire, Mike. 2000. "Policing by Risks and Targets: Some Dimensions and Implications of Intelligence-Led Crime Control." *Policing and Society* 9 (4): 315–36.
- McCulloch, Jude, and Dean Wilson. 2016. *Pre-Crime: Pre-Emption, Precaution and the Future*. London: Routledge.
- Nhan, Johnny. 2013. "Police Culture." In *The Encyclopedia of Criminology and Criminal Justice*, edited by Jay S Albanese, 1–6. Oxford, UK: Blackwell Publishing.
- Pasquale, Frank. 2015. *The Black Box Society: The Secret Algorithms that Control Money and Information*. Cambridge, MA: Harvard University Press.
- Pearsall, Beth. 2010. "Predictive Policing: The Future of Law Enforcement?" *NIJ Journal* 266: 16–19.
- Perry, Walt L., Brian McInnis, Carter C. Price, Susan C. Smith, and John S. Hollywood. 2013. *Predictive Policing: The Role of Crime Forecasting in Law Enforcement Operations*. Santa Monica, CA: RAND Corporation.
- Pinch, Trevor J., and Wiebe E. Bijker. 1984. "The Social Construction of Facts and Artefacts: Or How the Sociology of Science and the Sociology of Technology Might Benefit Each Other." *Social Studies of Science* 14 (3): 399–441.
- Ratcliffe, Jerry. 2008. *Intelligence-Led Policing*. Cullompton, UK: Willan Publishing.
- Rich, Michael L. 2016. "Machine Learning, Automated Suspicion Algorithms, and the Fourth Amendment." *University of Pennsylvania Law Review* 164 (4): 871–929.
- Risse, Mathias, and Richard Zeckhauser. 2004. "Racial Profiling." *Philosophy and Public Affairs* 32 (2): 131–70.
- Robinson, David, and Logan Koepke. 2016. *Stuck in a Pattern: Early Evidence on ‘Predictive Policing’ and Civil Rights*. Upturn. www.upturn.org/
- Rosenblat, Alex, Tamara Kneese, and danah boyd. 2014. "Predicting Human Behavior." Data & Society Research Institute. <https://datasociety.net>.
- Schauer, Frederick F. 2003. *Profiles, Probabilities, and Stereotypes*. Cambridge, MA: Belknap Press of Harvard University Press.
- Schauer, Frederick F. 2018. "Statistical (and Non-Statistical) Discrimination." In *The Routledge Handbook of the Ethics of Discrimination*, edited by Kasper Lippert-Rasmussen, 42–53. London: Routledge.
- Schmidt, Christine. 2018. "Holding Algorithms (and the People Behind Them) Accountable Is Still Tricky, but Doable." Nieman Lab, March 21, 2018. www.niemanlab.org.
- Selbst, Andrew D. 2017. "Disparate Impact in Big Data Policing." *Georgia Law Review* 52 (1): 109–95.

- Shapiro, Aaron. 2017. "Reform Predictive Policing." *Nature News* 541 (7638): 458–60.
- Sheehey, Bonnie. 2019. "Algorithmic Paranoia: The Temporal Governmentality of Predictive Policing." *Ethics and Information Technology* 21 (1): 49–58.
- Stroud, Matt. 2014. "The Minority Report: Chicago's New Police Computer Predicts Crimes, but Is It Racist?" *The Verge*, February 19, 2014. www.theverge.com.
- Tayebi, Mohammad, and Uwe Glässer. 2016. *Social Network Analysis in Predictive Policing*. New York: Springer.
- Turvey, Brent E. 2012. *Criminal Profiling: An Introduction to Behavioral Evidence Analysis*, 4th ed. Burlington, MA: Academic Press.
- Uchida, Craig. 2005. "The Development of the American Police: An Historical Overview." In *Critical Issues in Policing: Contemporary Readings*, 5th ed., edited by Roger G. Dunham and Geoffrey P. Alpert. Long Grove, IL: Waveland Press.
- Walsh, William F. 2005. "Compstat: An Analysis of an Emerging Police Managerial Program." In *Critical Issues in Policing: Contemporary Readings*, 5th ed., edited by Roger G. Dunham and Geoffrey P. Alpert. Long Grove, IL: Waveland Press.
- Walsh, William F., and Gennaro F. Vito. 2004. "The Meaning of Compstat: Analysis and Response." *Journal of Contemporary Criminal Justice* 20 (1): 51–69.
- Willis, James J., Stephen D. Mastrofski, and David Weisburd. 2004. *Compstat in Practice: An In-Depth Analysis of Three Cities*. Washington, DC: Police Foundation.
- Willis, James J., Stephen D. Mastrofski, and David Weisburd. 2007. "Making Sense of COMPSTAT: A Theory-Based Analysis of Organizational Change in Three Police Departments." *Law & Society Review* 41 (1): 147–88.
- Winner, Langdon. 1980. "Do Artifacts Have Politics?" *Daedalus* 109 (1): 121–36.
- Zarsky, Tal. 2013. "Transparent Predictions." *University of Illinois Law Review* 2013 (4): 1503–69.