BIG DATA POLICING: EVIDENCE AND APPEAL

DRAFT

Kevin D. Haggerty

Introduction

For years scholars have identified the role that computers and massive databases play in the dynamics of contemporary surveillance (Ball, Haggerty, and Lyon 2012). Indeed, Gary Marx was characterizing this as the ‘new surveillance’ close to three decades ago (Marx 1988). In the past few years, however, ‘big data’ has emerged as a type of faith that important but unpredictable gains in organizational efficiency, profit, and assorted other desirable institutional outcomes lie undiscovered in the mountains of accumulated computer data; if only we can develop the algorithms to identify them (Mayer-Schönberger and Cukier 2013).

Big data revolutions are anticipated in healthcare, transportation, and commerce. My comments here are confined to the realm of policing and security, a context where we are seeing both new big data initiatives and a corresponding rise in critical reflections on the desirability and legality of such developments (McCulloch and Wilson 2016). Here ‘policing’ is understood as the ambit of a diverse range of agencies and actors working to secure particular people, domains and processes; not exclusively the public police.

In these brief comments I do three things. First, I give a flavor for some of these initiatives in policing, and the logic of how they are meant to work. Second, I outline some of the difficulties in trying to assess their effectiveness. The concluding discussion points to one important reason why police officials may be drawn to big data policing strategies, irrespective of whether they 'work' in some instrumental fashion.

Technological and Informational Heterogeneity

The general picture of the emergence of big data policing initiatives is one where the greater availability of computerized information is changing both the focus and temporal horizon of policing. More specifically, the focus is increasingly on a series of categories and profiles derived from the analysis of information, as discussed below. Second, this has pushed officials from a type of reactive model of policing, to one that attempts to target in advance groups that automated analysis of datasets revels to be particularly dangerous, risky, or problematic. This is the much discussed 'pre-crime' component of big data policing. This temporal distinction can be too starkly drawn, as much of policing has always had a forward looking face, but there is undeniably a greater appetite for relying on data to prospectively select targets for intervention and protection.

In the weeks immediately after the 9/11 terrorist attacks John Poindexter proposed that the newly formed Department of Homeland Security create a new body to serve as a centralized data warehouse, collecting and storing any and all information about American citizens That information would then be made routinely available to police and security agencies. This initiative went by the Orwellian moniker of 'Total Information Awareness' (TIA). Even in those emotionally charged days, however, this level of sweeping centralized governmental surveillance capacity spooked both the American public and legislators. Poindexter's plan was quickly shelved.

In the ensuing years security officials have continued to yearn for as much information about the citizenry as possible. But rather than produce a massive centralized state data collection agency, what has emerged is instead what might be best characterized as a neoliberal surveillance complex. This is characterized by efforts to capitalize on the de-facto centralizing possibilities inherent in the ability to integrate the diverse data collection efforts already underway by assorted state and private agencies without the need for a new governmental bureaucracy. The aim is to bring together, analyze, and make actionable the information stored in widely dispersed public and databases. As Harcourt notes, this means that what is now characteristic about state surveillance is that it is not confined to 'the state,' but instead involves "an amalgam of various national intelligence services, Google, Microsoft, other Silicon Valley firms, Facebook and other social media corporations, private surveillance industry companies and consultants, IT departments everywhere, and…. Local police departments, friends, hackers, and curious interlopers" (Harcourt 2015: 72).

Indeed, a characteristic attribute of much big data policing is that it displays a type of assemblage methodology (Haggerty and Ericson 2000), that aims to combine and integrate a heterogeneous range of previously discrete actors, agencies, technologies, and data into an operational whole in hopes of providing a more refined or operationally useful understanding of places, people, and processes. Officials are drawn to this style of neoliberal surveillant assemblage for several reasons. It fits with the contemporary political ethos critical of all governmental bureaucracies for being expansive and inefficient as compared to private firms. State officials recognize that it is a comparatively easy and effective way to access widely dispersed information. In some cases is also allows security officials to engage in a type of surveillance legal gerrymandering, whereby the police get access to personal information that they are explicitly precluded from collecting themselves.

One consequence of this heterogeneity is that it is hard to fully assess how (and if) big data is transforming policing. The real action surrounding big data policing is taking place amongst a widely dispersed and relatively independent set of planers, statisticians, technology start ups; all of whom are working away to develop their own proprietary software. So neoliberal big data policing does not have a key initiative or main bureaucracy that we can easily study. Instead, there is a complex mass of technological initiates existing in assorted states of planning, development, promotion, testing, and operation. These initiatives employ diverse methodologies, draw on different data sets, and make quite different claims to success - or even how ‘success’ might be conceptualized. Hence contemporary discussions of policing and big data tend too easily move across mention of currently operating systems, prototypes, and speculative fantasies.

One of the earliest and most celebrated big data policing initiatives was the COMPSTAT system pioneered by the New York Police Department. COMPSSTAT was one part police accountably mechanism, one part resource allocation structure, and several parts police theatre—with police officials being publically and ceremonially abused for failing to reduce crime rates. What distinguished COMPSTAT from subsequent big data policing initiatives was that COMPSTAT relied primarily on official crime and arrest data, and was not particularly sophisticated in terms of the analysis of those data. Over time those crime data were augmented by information drawn from other sources in efforts to predict crime patterns, including information about the weather, season, and time of the week—giving a hint of what was to come in big data policing. In the ensuing years initiatives to track the geolocation of crime trends have become more sophisticated and automated through programs such as BlueCrush. Currently used by a Memphis police department BlueCrush uses statistical modeling of past crime data to identify crime “hot spots.” Police are then directed to those hot spots to try and deter crime and make arrests (Miller 2014: 117).

A more recent initiative gives a sense of the diverse data sources that software companies are now drawing upon. *Beware,* produced by the company *Intrado,* is being used by the Fresno California police department. As an officer in Fresno s dispatched to an event, the address or name of the individual(s) involved are quickly searched on the *Beware* system, producing a color coded (green, yellow, red) threat level. These threat levels are derived from a computerized process that combines standard police record system with information that historically has not been officially been part of a person’s criminal profile, including vehicle registrations, property records, Facebook posts, tweets, etc. (Jouvenal 2016).

A comparable logic is apparent in the pre-flight screening system conducted by the Transportation Security Agency in the United States. This system scans travellers prior to their arriving at the airport. Their passport details are searched against a series of government law enforcement databases, but also against databases maintained by the Internal Revenue Service as well as airline frequent flyer programs, and credit risk scoring agencies. Again, the aim is to slot travellers into ‘risk levels,’ with individuals deemed to be higher on the risk hierarchy receiving a greater degree of subsequent scrutiny when they arrive at the airport (Miller 2014: 116).

A more futuristic example of efforts to use diverse data points to identify risky individuals is the Future Attribute Screening Technology (FAST), being developed by the Department of Homeland Security. The designers anticipate using this system in airports or other potential targets for terrorism. The technology itself involves integrating a series of sensors, video cameras, and audio recorders to remotely collect behavioral data about a person as they move through a crowd. The technologies that have been identified as part of this system would collect information on pheromones, skin conductivity, eye blink rate, repertory patterns, skin conductivity, all without the target's knowledge. The aim is to use this information to identify individuals whose physiology is elevated because they plan on engaging in some kind of crime or violence.

Evaluation

Do such strategies work? It is a recurrent question, and one that is actually more difficult to assess than is typically assumed. While one often hears claims of successes – sometimes dramatic successes - closer attention to a series of methodological issues familiar from the longstanding criminological effort to evaluate anti-crime initiatives gives reason for pause.

Part of the immediate problem is that much of what we know about big data systems is anecdotal. For the many systems that are currently in production we have little beyond vague promises about their revolutionary potential. For existing initiatives, the claims to success have typically been voiced by the people who have a financial interest in the program, or by the officials who have staked their reputations on implementing such programs. As such, many of the 'successes' more closely resemble promotional copy than the findings of rigorous and independent social scientific evaluation.

An additional complicating factor in determining what does and does not work concerns the dynamics of state and corporate secrecy. Although big data is increasingly used in the realm of national security to single out particular individuals or groups for greater police attention, state officials are predictably reluctant to provide the details of how these systems work, and the specifics of how their successes are measured. And while security concerns are legitimate, the attendant need for secrecy makes it extremely difficult to independently assess whether these systems work, how well they work, and how cost effective they might be – particularly in comparison to other potential policing strategies. The end result is often that state officials explicitly or implicitly tell citizens 'trust us, this is working.' Unfortunately, time and again when citizens ultimately get access to information on the real world of security, they have often discovered that secretive policing agencies have violated such trust, implementing programs that were ineffective, politically partisan, ideologically blinkered, wasteful, unconstitutional, or blatantly illegal (Rosenfeld 2012). Most recently the Snowden revelations have given the public a glimpse of the extent to which security agencies are willing to bend or break the law in their attempts to access and use personal data (Lyon 2014, Greenwald 2014).

Such situations make it difficult to identify and assess the successes of big data national security initiatives. For example, the NSA's much discussed eavesdropping program uses metadata and big data analytics to ostensibly identify possible terrorist targets and thwart terrorist attacks. The Obama administration at one point indicated that such eavesdropping had allowed the NSA to thwart fifty four terrorist plots. When subsequently pressed for details, the NSA Director admitted that he could only point to one instance of demonstrable success – a case in which a Somali immigrant who had donated money to al-Shabaab (Miller 2014: 118). Even in that case, however, it is not clear that identifying a person who financially aided an established terrorist organization can easily be characterized as 'thwarting a terrorist plot.'

Private corporations are also often not forthcoming on the specifics of how their big data programs work, although the operational logic here is focused on protecting trade secrets. Consider, for example, the *Beware* program mentioned earlier, which culls masses of information to produce a risk category for individuals. To understand how risk scores are arrived at, one need access o the algorithms that both amass all the data, and vitally, which help to make disperse bits of information comprehensible and actionable – in this case in the form of a risk score (Amoore and Piotukh 2015). Those algorithms determine what information is and is not included, over what timeline, within what geographical parameters, and crucially, how those dispersed bits of information are differentially weighted to produce a risk score. Understanding the specifics of what information is collected and how it is weighted is key to evaluating the successes of these systems, and the extent to which they might discriminate against particular groups, or unfairly compound the social disadvantages such individuals face. Unfortunately, the corporation that created *Beware* will not reveal their algorithms, as they see them as part of their trade secrets. This emphasis on trade secrets is a common problem limiting our ability to assess the operation of big data policing systems.

James William's (Williams 2009, 2013) research into some aspects of financial big data surveillance gives a rare glimpse into the extent to which the construction of algorithms, in combination with legal regimes, and differential access to data, can shape the ostensible successes of big data initiatives. As Williams details, the Toronto Stock Exchange contracts with the private organization Investment Industry Regulation Organization of Canada (IIROC) to identify forms of market manipulation, improper trade execution, and front-running. IIROC is then given access to the massive amounts of transactional data produced by the Toronto Stock Exchange.

The volume of data available to regulators means that they face serious challenges in making sense of an almost unmanageable data glut. The solution which is now familiar across big data policing initiatives has been for investigators to design computer algorithms to identify suspicious activities. These algorithms monitor the movement of funds between accounts, and the purchase and sale of stocks, bonds and other financial instruments, and automatically trigger an alert when they identify suspicious activity. That might happen, for example, when a trading rule has been violated – say a designated company 'insider' has traded some of his or her stocks. Alternatively, an alert might sound when trading occurs outside of established norms for a stock's price or for the volume of that stock traded in the recent past. In essence, these algorithms monitor for dramatic spikes in stock prices and trading volumes that occur outside of a statistically determined normal range. These statistical norms are themselves based on one month rolling averages. So, if a stock that last week traded at comparatively leisurely rate starts to sell rapidly, the computer produces an alert. Thousands of such alerts can sound in a single day, which the IIROC staff then try and decipher to determine whether they are explainable aberrations, or if they warrant further investigation.

The regulatory understanding about the markets depends almost entirely on what data are scrutinized, and the design of the algorithms used to make sense of that information. The result is a series of notable regulatory blind spots. For example, the fact that the alert system only monitors a small number of market indictors, including volume, price, and volatility, provides a comparatively thin and rudimentary view of markets. Second, jurisdictional issues mean that the IIROC regulators only focus their attention on equities-based markets and cannot monitor activities on the bond, derivatives, and commodities markets. Consequently, they cannot detect the common practice of cross-market manipulation. Finally, the fact that the alert system is based on detecting statistical variations based on one month averages means that financial actors involved in more long term market manipulations are invisible to the regulations. All of this leads Williams (2009: 481) to conclude that "[w]hile it may be true that digitized financial markets are more transparent than ever, they are rendered surprisingly opaque by virtue of the superficiality of these representations and the surfeit of information relative to the paucity of ‘true’ market knowledge.” Moreover, it is the established powerful actors who are best positioned to take advantages of these regulatory blind spots.

The precise details how algorithms are constructed is therefore key to giving the authorities a particular vision of risky people, places and processes. And while big data proponents are enthusiastic about the possibilities inherent in the ability to collect massive amounts of information, the sheer scope of this information as compared to comparatively small number of targets to be identified can make interventions programmatically difficult. In a much discussed article Bruce Schneier (Schneier 2006) details some of the pragmatic difficulties in trying to identify terrorists through analysis of risk profiles, focusing on the difficulties inherent in trying to strike a manageable balance between Type 1 and type 2 errors. In essence, officials relying on scrutinizing massive data trails to identify a small group of terrorists are faced with the dilemma of either identifying far too many innocent individuals in order to catch a large proportion (but still not all) of terrorist suspects, or to calibrate the identification system such that fewer innocent individuals are stopped, but which will then result in an even smaller percentage of legitimate terrorist suspects being identified. In both scenarios, the system would be incapable of identifying all of the risky individuals.

For programs that are already being used by the police, the question becomes focused on how success is to be measured. A key difficulty in such evaluations is to try and unpack the influence of a particular initiative, as compared to the influence of a host of other potentially relevant factors. Crime control is a constantly shifting and evolving field, where unlike in a laboratory, it is difficult to control for “all other potentially relevant factors.” So a big data initiative might be implement at the same time that a dramatic crime is being extensively covered in the media, new police reporting mechanisms are being introduced, a larger cohort of new police recruits are entering street, a new illegal narcotic is circulating within particular population groups, and so on. Any of these things phenomena either individually or in combination can impact the crime or arrest rate independent of any influence the big data imitative might have had. It is notoriously difficult to unpack the differential influences of such diverse factors.

One gets a flavor of this from the discussions about the purported successes of the COMPSTAT system in New York. While police officials pointed to a much celebrated decline in violent crime which they attributed to COMPSTAT, subsequent analysts have noted that the New York crime drop is in line with a major decline in crime rates that occurred in the United States at that time, irrespective of whether those other police forces were employing a COMPSTAT system. The implication being that the successes claimed by COMPSTAT can actually be attributed to a wider set of societal or demographic factors that were apparent across the country.

Unfortunately, most studies of big data do not take into account the prospect that changes in crime or arrests rates are due to multi-dimensional and interactional factors. Instead, they typically rely on a simple before and after methodology, where crime or arrest rates are measured before and after an intervention, with any fluctuations in crime or arrest trends being attributed to the intervention. This ‘before and after’ methodology is typically seen as the weakest evaluative approach, precisely because it ignores the prospect that other individual or interactional factors might have played a role in producing any demonstrated changes (Sherman et al. 2006).

It can also be unclear what statistical measures should be taken as evidence of success. So, police officials might point to a decline in the crime rate in a particular area as evidence of the success of a big-data initiative. This typically overlooks whether the crime has simply been displaced to a different area – something that routinely happens with particular types of crime, such as prostitution and drug dealing.

Alternatively, sometimes the police point to increase arrest rates as evidence that an anti-crime initiative is working. One sees this, for example, in relation to the previously mentioned BlueCRUSH program that identifies ‘hot spots’ for increased police intervention. The increased arrest statistics that result from such intensified policing are then used to point to the successes of the initiative. However, such an approach ignores the criminological truism that if you intensively monitor *any* population group, you are going to find more criminal behavior. The specific types of crime and level of crime will be different for different population groups, but it is a circular logic and self-fulfilling prophecy to intensively police a particular group and then use the fact that you made increased arrests amongst members of that group as a way to justify that increased police attention.

All of these provide just a taste of the many difficulties inherent in program evaluation that a long tradition of evaluative research in criminology has identified. When taking all such factors into account is apparent that the successes of big data policing remain inconclusive. That is not to say that big data strategies do not work, but that as often (perhaps routinely) occurs, the roll out of new information technology in policing occurs in an environment of ambiguous and contradictory evidence as to their successes. For those individuals who believe in the promise of evidence based policing, such an ambiguous evidentiary situation might be seen as a reason to ‘go slow’ on the development and introduction of big data initiatives. Yet the exact opposite seems to be occurring,

Discussion: The Unarticulated Appeals of Big Data Policing

It is easy to understand why the police might be drawn to the promise that big data will reduce crime and victimization rates, even if the evidence for such successes remains inchoate. It is also the case that other factors pertaining to institutional self-interest, budgets, and broader political dynamics and cultural currents, can encourage organizations to embrace new technologies. For example, for the police the appeal of big data can be subtly informed by dominant understandings of modernity, with progress in the contemporary West being regularly associated with using information technology. The police can be attracted to information technology—including big data—because of a forward looking desire to appear to be on the cutting edge of policing developments, or alternatively due to not wanting to be seen as being left behind by modern approaches to policing.

By way of concluding and admittedly speculative comment, I want to suggest that part of the latent and perhaps unacknowledged appeal of big data strategies for the police lies in how such approaches help to address a vexatious political dilemma. In particular, big data policing helps the police avoid accusations of bias in a political climate where concerns about racism and discrimination have become inescapable in some jurisdictions.

For decades the police, particularly in the United States, have been accused of exercising their discretion in racist ways. Critics of the police—as well as a series of official inquires and investigations—have pointed out that the discretion of both individual line officers and police managers, has culminated in significantly higher numbers of racialized individuals being subject to stop and search, being pulled over for 'driving while black,' and being charged with minor offences. Some see intensive policing strategies designed to 'crack down' on certain types of crime as akin to an undeclared war on racialized urban men. All of this has culminated in the dramatic over-representation of black and Latino individuals in American prisons (Alexander 2010). In the eyes of a good many individuals, police decisions about who to stop, arrest, and what neighborhoods to target, have become synonymous with an implicit, and sometimes explicit, police racism.

It is here that I believe big data policing initiatives have a latent appeal for police officials. Big data gives many types of police decisions a veneer of objective scientism, helping to shelter individual officers and police organizations from accusations of bias. This insight follows from Ted Porter's (Porter 1994, 1995) study of the army corps of engineers, which documented the historical rise of a 'trust in numbers.' As Porter's shows, decisions about the location of bridges, dams, and roadways traditionally relied on the professional judgment of engineers. However, as engineers faced increasing accusations that their decisions were shaped by politics, their profession embraced a series of quantified and standardized accounting rules to help guide/determine the scope and location of engineering projects. For Porter, the problem that such numerical accounting rules solved was not about where to build bridges, but how to insulate a profession from accusations that its members were making politically partisan decisions.

I suspect that part of the appeal of big data for the police is roughly comparable. By using big data analytics to make decisions about what crime 'hot spots' need intensive policing, the police partially insulate themselves from accusations that such decisions are biased along racial or class lines. So, police officials can, for example, claim that the intensive police crackdowns in certain neighborhoods are not the result of some subjective and potentially suspect police decision making process, but are instead derived from a form of mechanical objectivity based on big data analytics. This is notwithstanding the fact that I suspect that if individual police officers were polled about what neighborhoods need extra policing due to gang violence, for example, that their subjective responses would probably not differ much form those suggested by the big data analytic. The same is likely true of the list of ‘top 100’ criminals that big data programs suggest deserve increased police scrutiny. Long before big data the police worked on the assumption that there was a subset of particularly notorious individuals in society who needed extra police attention.

Officers who pull over individuals arriving at the airport because ‘the computer told them’ to do so, can rest comfortable knowing that they will not be accused of racism or bias. Likewise, officers criticized for taking a particularly aggressive stance when dealing with a citizen can justify their behavior by pointing to the fact that an algorithm told them that this individual might be ‘high risk.’ None of this is to say that subjective bias and racism disappear from police decision-making, but it is instead displaced from the discretion of individuals and instead becomes embedded into the inscrutable and often secret algorithms that are seen as the wave of the future in policing.

References

Alexander, Michelle. 2010. *The New Jim Crow: Mass Incarceration in the Age of Colorblindness*. New York: The New Press.

Amoore, Louise, and Volha Piotukh. 2015. "Life Beyond Big Data: Governing With Little Analytics." *Economy and Society* 44 (3):341-366.

Ball, Kirstie, Kevin Haggerty, and David Lyon. 2012. *The Routledge Handbook of Surveillance Studies*. London: Routledge.

Greenwald, Glenn. 2014. *No Place to Hide: Edward Snowden, the NSA, and the U.S. Surveillance State*. New York: Random House.

Haggerty, Kevin D., and Richard V. Ericson. 2000. "The Surveillant Assemblage." *British Journal of Sociology* 51 (4):605-22.

Harcourt, Bernard E. 2015. *Exposed: Desire and Disobedience in the Digital Age*. Cambridge: Harvard University Press.

Jouvenal, Justin. 2016. "The New WayPolice Are Surveilling You: Calculating Your Threat 'Score'." *The Washington Post*, January 10. https://www.washingtonpost.com/local/public-safety/the-new-way-police-are-surveilling-you-calculating-your-threat-score/2016/01/10/e42bccac-8e15-11e5-baf4-bdf37355da0c_story.html.

Lyon, David. 2014. "Surveillance, Snowden, and Big Data: Capacities, Consequences, Critique." *Big Data & Society* 1:1-13.

Marx, Gary. 1988. "The New Surveillance." In *Undercover: Police Surveillance in America*. Berkeley: University of California Press.

Mayer-Schönberger, Viktor, and Kenneth Cukier. 2013. *Big Data: The Revolution That Will Transform How We Live, Work, and Think*. Boston: Houghton Mifflin Harcourt.

McCulloch, Jude, and Dean Wilson. 2016. *Pre-Crime: Pre-emption, Precaution and the Future*. London: Routledge.

Miller, Kevin. 2014. "Total Surveillance, Big Data, and Predictive Crime Technology: Provacy's Perfect Storm." *Journal of Technology, Law and Policy* 19:105-146.

Porter, Theodore. 1994. "Objectivity as Standardization: The Rhetoric of Impersonality in Measurement, Statistics, and Cost-Benefit Analysis." In *Rethinking Objectivity*, edited by A. Megill, 197-237. Durham: Duke University Press.

Porter, Theodore. 1995. *Trust in Numbers: The Pursuit of Objectivity in Science and Public Life*. Princeton: Princeton University Press.

Rosenfeld, Seth. 2012. *Subversives: The FBI's Sar on Student Radicals, and Reagan's Rise to Power*. New York: Picador.

Schneier, Bruce. 2006. "Why Data Mining Won't Stop Terror." *Wired*, March 9.

Sherman, Lawrence W., David P. Farrington, Brandon C. Welsh, and Doris Layton MacKenzie, eds. 2006. *Evidence-Based Crime Prevention: Revised Edition*. London: Routledge.

Williams, James. 2009. "Envisioning Financial Disorder: Financial Surveillance and the Securities Industry." *Economy and Society* 38 (3):460-491.

Williams, James. 2013. "Regulatory Technologies, Risky Subjects, and Financial Boundaries: Governing 'Fraud' in the Financial Markets." *Accounting, Organizations and Society* 38:544-558.