

Ethical Theory and Moral Practice

The Relative Moral Risks of Untargeted and Targeted Surveillance

--Manuscript Draft--

| | |
|--|--|
| Manuscript Number: | |
| Full Title: | The Relative Moral Risks of Untargeted and Targeted Surveillance |
| Article Type: | Original Research |
| Keywords: | criminal justice; surveillance; privacy; stigmatisation; discrimination |
| Corresponding Author: | Katerina Hadjimatheou, Ph.D. University of Birmingham Birmingham, UNITED KINGDOM |
| Corresponding Author Secondary Information: | |
| Corresponding Author's Institution: | University of Birmingham |
| Corresponding Author's Secondary Institution: | |
| First Author: | Katerina Hadjimatheou, Ph.D. |
| First Author Secondary Information: | |
| Order of Authors: | Katerina Hadjimatheou, Ph.D. |
| Order of Authors Secondary Information: | |
| Abstract: | Is surveillance that is targeted towards specific individuals easier to justify than surveillance that targets broad categories of people? Untargeted surveillance is routinely accused of treating innocent people as suspects in ways that are unfair and of failing to pursue security effectively. I argue that in a wide range of cases untargeted surveillance treats people less like suspects than more targeted alternatives. I also argue that it often deters unwanted behaviour more effectively than targeted alternatives, including profiling. In practice, untargeted surveillance is likely to be least costly morally and most efficient as a means of enforcing the rules of a specific activity or institution. Targeted alternatives are likely to be more appropriate means of enforcing the rules of criminal law in general. |

Title: **The Relative Moral Risks of Untargeted and Targeted Surveillance**

Author: Katerina Hadjimatheou

Biographical Note: Katerina Hadjimatheou is a Research Fellow at the Centre for the Study of Global Ethics in the Department of Philosophy at the University of Birmingham, UK. She received her PhD from the University of Essex in 2010. She is currently working on an EU-funded project on the ethics of surveillance.

Contact Details:

Katerina Hadjimatheou

The Centre for the Study of Global Ethics

The Department of Philosophy

The University of Birmingham

Birmingham

United Kingdom

B15 2T

hadjimak@bham.ac.uk

01214145600

07837549931

The Relative Moral Risks of Untargeted and Targeted Surveillance

Abstract: Is surveillance that is targeted towards specific individuals easier to justify than surveillance that targets broad categories of people? Untargeted surveillance is routinely accused of treating innocent people as suspects in ways that are unfair and of failing to pursue security effectively. I argue that in a wide range of cases untargeted surveillance treats people less like suspects than more targeted alternatives. I also argue that it often deters unwanted behaviour more effectively than targeted alternatives, including profiling. In practice, untargeted surveillance is likely to be least costly morally and most efficient as a means of enforcing the rules of a specific activity or institution. Targeted alternatives are likely to be more appropriate means of enforcing the rules of criminal law in general.

Keywords: Surveillance; privacy; stigmatisation; discrimination; reciprocity

Introduction

Untargeted surveillance is the surveillance of all individuals in a specific place or engaged in a specific activity. It is used for a wide range of purposes: to improve productivity in the workplace by monitoring employee computer or lavatory usage, to ensure fairness in sports competitions by conducting random drug tests, to enforce speeding laws on roads, and to enforce the rules of gambling in casinos. It is used in schools, where spot-checks of lockers for drugs, alcohol or weapons sometimes take place, and hospitals, where CCTV is often installed.

Untargeted surveillance is routinely used by police and other state security officials. Examples include the use of metal detectors at airports, fingerprinting and other biometric recordings for identity documents, searches or frisks on entrance into government buildings or prisons, stops and searches at transport hubs and in city centres, and Criminal Records Bureau (CRB) checks for those applying to work with children or other vulnerable people. Less common examples include in-depth investigations by security services into the lives of people taking up positions with high-level security clearance in the government or intelligence services.

Surveillance can be more or less targeted, and what I have been calling untargeted surveillance occupies one end of the spectrum. At the opposite end lie measures of surveillance that act on specific information linking particular individuals to particular incidents of unwanted behaviour (what in legal circles is called 'individual' or 'individualised suspicion' [Clancy, 1994]). At different points between these two extremes lie measures such as random screening, 'fishing expeditions', and profiling.

Untargeted surveillance is on the rise. This is partly because of technological developments that make it both easier and less costly to implement. Thus software now enables

employers to monitor employee computer use; metal detectors and body scanners allow fast surveillance of all air passengers; an increasing variety of information can be stored on databases that are easily searchable; and scanning and testing technology enables fast and relatively un-invasive collection of biometric data and drug testing. There is also increasing emphasis by governments, police forces, and other agencies on the effective prevention of hazardous, wrongful, and criminal behaviour. Untargeted surveillance is primarily a preventive approach to reducing security threats and rule-breaking in general.

This paper considers and rejects two commonly voiced concerns about the fairness and effectiveness of untargeted surveillance. These are summarised below.

1. *Untargeted surveillance is unfair because it treats innocent people like suspects without any prior evidence of suspiciousness.* Suspicion should be triggered by evidence of wrongdoing, rather than precede it. By suspecting those for whom no evidence of wrongdoing exists, untargeted surveillance stigmatises them unfairly; undermines the presumption of innocence (Haggerty and Ericson 1997: 42; Monahan, 2010: 99); and “promotes the view...that everybody is untrustworthy” (Norris in House of Lords, 2009: 27[107]; New Scientist, 2006).
2. *Untargeted surveillance is inefficient in reducing security threats.* Evidence about the effectiveness of untargeted surveillance strongly suggests that such measures result in the detection of very few if any genuine security threats of the kind it is intended to reduce.¹ Scrutinising individuals who are obviously not suspicious imposes privacy and other costs on them without any prospect of increased benefit to security and is thus “futile and time-wasting” (Lord Bingham, UKHL 12, 2006, 35). Targeted surveillance focuses scarce resources only on those for whom suspicion exists and is therefore likely to be more efficient (Bou-Habib, 2008; Risse and Zeckhauser, 2004).

Each of these criticisms implies that the risks of untargeted surveillance could be reduced or eliminated by imposing on those doing the surveillance a prior requirement for objective, evidence-based, targeted suspicion.

I argue that, on the contrary, in a wide range of cases untargeted surveillance either produces none of the negative outcomes identified in 1 above or carries an equal or lower risk of producing them than targeted alternatives. I also argue that untargeted surveillance can be more efficient a means of deterring unwanted behaviour than targeted approaches. I put forward a preliminary model for understanding the relative risks and benefits of targeted and untargeted surveillance. According to this model, the more targeted the surveillance is, the higher the risks of stigmatisation, the more intrusive the privacy

¹ “As well as the potential counterproductive effect of stop and search under section 44, it has also not proven to be an effective tool in countering the terrorist threat. Statistics demonstrate that as little as 0.6% of stop and searches under section 44 in 2008/9 resulted in an arrest.” (Liberty 2010:54). Similar views are expressed in Human Rights Watch (2010:48-9) and Moeckli (2007: 668).

interference, the fewer individuals affected and the more efficient it will be as a detector of wrongdoers. The less targeted the surveillance is, the lower risk of stigmatisation, the less intrusive the privacy measure, the greater the number of people affected and the more efficient it is likely to be as a deterrent or wrongdoing. Partially targeted measures of surveillance such as profiling present the highest risk of discrimination. Untargeted surveillance is often easier to justify than targeted surveillance when it is used to enforce the rules of a specific activity or institution. It is less easy to justify as an approach to enforcing the rules of the criminal law in general.

1. Untargeted surveillance and stigmatisation

I begin this section by discussing some difficulties with the idea that untargeted surveillance treats people like suspects and thereby stigmatises them. Stigmatisation and its costs result from being singled out. Untargeted surveillance is, other things being equal, *less* stigmatising than targeted surveillance, precisely because it treats everybody within its range with equal scrutiny.

One standard understanding of stigmatisation defines it as the process of marking a person out as having an undesirable characteristic (Courtwright, 2011; Arneson, 2007). For the purposes of this paper, the undesirable characteristic associated with people who are marked out as suspects is failure to meet the justified moral standards of the (relevant) community. When surveillance marks people out as suspects, it marks them out as potentially having failed to uphold or having violated a rule they are *prima facie* obliged to follow. For example, the rules of transport security, of professional codes of ethics, and of the criminal law are rules people are normally obliged to follow either because they have explicitly consented to do so or because they enjoy the benefits that accrue from a situation in which the rules are followed collectively.

The harms of stigmatising people as suspicious can affect both individuals and society as a whole. On an individual level, being stigmatised as having failed to maintain the moral standards of the community can be humiliating. People are humiliated when they cannot prevent appearing to others in ways that are demeaning (Bou-Habib, 2011:44). If individuals who are humiliated as a result of their stigmatisation are in fact innocent of any rule-breaking, they may feel anger or indignation at what they perceive to be an unjust implication of wrongdoing.² This may create knock-on social costs, by eroding their trust in the surveilling authorities, which in turn may reduce their willingness to cooperate with those authorities and to support the enforcement of other rules (FRA 2010: 21). Stigmatisation can also make those affected feel alienated from others as they perceive that others' estimation of them has fallen. This can affect negatively people's self-confidence and their willingness and ability to connect and engage effectively with others. It can also lead people to become socially isolated if they react by withdrawing from contact with others

(Courtwright, 2011:2). At the same time, social isolation of those stigmatised may be imposed from outside, if other individuals avoid or reject them as a result of surveillance policy. When particular groups of people who share salient traits -such as religion or race- are stigmatised as suspicious, as has been seen to occur with semi-targeted measures of surveillance such as profiling, this may exacerbate existing prejudices against them (Kennedy, 1997; Schauer, 1994; Lever, 2004), or even create new ones (Rumann and O'Connor, 2003).

When stigmatisation based on ethnic grounds reflects ethnic prejudice it can be experienced as particularly humiliating both by those surveilled and by those who are not themselves surveilled but who share the traits that trigger the suspicion and for that reason feel implicated (Parmar, 2012: 9). As all this suggests, stigmatisation caused by surveillance can be harmful to individuals and can distort social relationships in ways that impede the pursuit of important social goals, including security and equality.

All of these potential costs are likely to increase along with the importance of the moral standard or rule one is suspected of having broken. Being stigmatised as a suspected cheat at cards in a casino is neither as humiliating nor as likely to lead to social isolation or discrimination as being stigmatised by police as a potential paedophile. The costs of stigmatisation are also likely to intensify the greater the implication of guilt conveyed in the measure of surveillance: a bag search is less stigmatising than a house search, which in turn is less stigmatising than being taken to the police station for questioning. The extent to which people are stigmatised may also be affected by the number and identity of any witnesses to the surveillance: other things being equal, the greater the number of people who witness or become aware of the stigmatisation and the more influential or important those people are, the more severe both the individual and social costs are likely to be.

There are good reasons to think that both the extent to which surveillance treats people like suspects and its stigmatisation costs increases the *more targeted* the measure of surveillance. As has already been established, stigmatisation occurs when individuals are marked out as suspicious. Being marked out implies being identified in some way that distinguishes one from other members of the wider community or the relevant group. Being pulled out of line for further search or questioning at an airport; being stopped and searched on a busy train platform while other passengers are left alone; having one's travel history, credit card, and other records searched before flying because one fits a profile of a potential terrorist- these are all examples of being singled out and thereby marked out for suspicion. They are all stigmatising, because they all imply that there is something suspicious about a person that justifies the intrusion.

In contrast, untargeted surveillance such as blanket screening at airports, spot screening of all school lockers for drugs, and the use of speed cameras neither single people out for scrutiny nor enact or convey a suspicion that those surveilled are more likely than others to be breaking the rules. Rather, everybody engaged in the relevant activity is subject to the

same measure of surveillance, indiscriminately and irrespective of any evidence suggesting particular suspiciousness. Such evidence may well emerge from the application of untargeted surveillance, and that evidence may then be used to justify singling people out for further, targeted surveillance. But untargeted surveillance itself affects all people within its range equally and thus stigmatises none in particular.

It might be objected that the act of choosing a range for any measure of untargeted surveillance itself involves suspicion and that therefore untargeted surveillance continues to be stigmatising. This seems wrong for many routine examples of untargeted surveillance. There are reasons why people travelling by aeroplane are subject to searches not imposed on those travelling by car and reasons why car drivers but not pilots are subject to surveillance by speed camera. In order to accept the claim that surveilling all car drivers by speed camera stigmatises them as suspected speeders or that surveilling all airline passengers stigmatises them as potential terrorists, we would have to accept that these reasons are always or often related to the suspiciousness of either the activity surveilled or the individuals engaged in it. This seems highly unlikely. Driving cars is a dangerous activity and thus should be regulated by enforceable rules such as speed limits. The *only* people who can break the particular rule of road safety embodied in a speed limit are car drivers. Surveilling anyone else but them would be perverse. Therefore it seems wrong to argue that the indiscriminate surveillance of all people engaged in a particular activity singles them out and thereby stigmatises them, *when the surveillance is used in order to enforce the rules of that particular activity*.

In contrast, it would be correct to argue that singling out and stigmatisation occurs when additional measures of surveillance -over and above blanket speed checks- is targeted at drivers with existing traffic offenses or points on their licenses. In such cases, evidence of a potential for rule-breaking is used as a reason to subject some drivers but not others to preventive intrusion. The same goes for a measure that surveilles all drivers in order to enforce a rule unrelated to driving, such as the rule against tax evasion. If driving a certain kind of car were a strong indicator of potential tax evasion, then singling out the relevant car drivers from all other taxpayers as targets of surveillance would stigmatise them.

However, even if this argument about the nature and causes of stigmatisation is correct, it does not address the apparent fact that those affected by untargeted surveillance might sometimes *feel* stigmatised. For example, the British Athletes Commission recently objected to the proposed use of random or blanket drug testing in the London 2012 Olympic Village on the ground that it would “make athletes feel like criminals” (BBC, Jan 4, 2010). Making athletes feel like criminals seems unfair, even if there are potential security benefits to doing so, because participation in competitive sports events is an innocuous activity and only a tiny minority of athletes actually break the rules. The same might be argued for many other activities surveilled indiscriminately.

Two things can be said about this. Firstly, if we accept the analysis of stigmatisation offered above, then the only stigmatisation costs that arise when people feel they are being stigmatised but are not actually being singled out are costs connected to those feelings. These would include feelings of humiliation and indignation and the loss of trust in authorities such feelings might provoke. But they would not include discrimination or exclusion by others. Authorities can reduce the extent to which stigmatisation-related feelings result from untargeted surveillance by explaining to those surveilled that they are not suspected of any wrongdoing or crime, but that the surveillance is a preventive measure which can be shown to be both necessary and proportionate to the enforcement of rules protecting important collective interests. If such precautions are taken, then it is unlikely that genuinely untargeted surveillance will provoke stigmatisation-related feelings of sufficient vehemence to significantly undermine trust between individuals, surveilling authorities, and the community as a whole.

However, it may never be possible to eliminate entirely the negative feelings triggered in people as a result of their encounters with surveilling authorities. Measures that intrude on one's privacy can provoke feelings of indignation, resentment and humiliation even when the policy that leads to their imposition is just and even when one both consents to and recognises the justice of that policy. Having one's bags rummaged through or one's intentions questioned are not pleasant experiences. Authorities should recognise this and should take action to make surveillance as un-invasive as possible and to ensure that the encounter with the surveilling authority is civil and does not involve abuse. Any negative feelings that remain after such efforts have been made may still count among the negative consequences of surveillance. But the fact that these feelings are based in something other than a reasonable belief that one is being stigmatised or treated unfairly should affect the weight they are accorded in any moral assessment of the measure that triggered them (Risse and Zeckhauser, 2004:149; Reiman, 2010:15). Other things being equal, the less reasonable the emotional reaction, the less moral weight it should receive.

Untargeted surveillance should be distinguished from profiling, with which it often has much in common, but which does treat people like suspects and thereby stigmatise them. Profiling can be defined as the systematic association of sets of physical, behavioural, psychological or ethnic characteristics with rule-breaking and their use as a basis for making security decisions. It is increasingly proposed as an alternative to untargeted surveillance, despite the widespread concerns about its moral risks.³ Profiling is similar to untargeted surveillance to the extent that it affects *groups* of people, but it differs in that it singles these groups out for suspicion. For example, the use of metal detectors or sniffer dogs in some 'problem' schools resembles untargeted surveillance because it is used to enforce the

³ For instance, the latest EU Security Research call invites proposals for security checks "based on threat levels and a dynamic evaluation of risks at individual level, instead of the current [100% checks] scheme". 7.Topic SEC-2013.3.4-3 Security checks versus risk at borders – Capability Project. EC FP7 Security Work Programme 2013.

rules of a particular institution and is targeted indiscriminately at all members of the institution who could break the rules. But the fact that such measures respond to suspicion that pupils at some schools are more likely to be breaking the rules than others makes them more like profiling.

Many of the most common objections to profiling centre on its apparent failure to fulfil a moral requirement of reciprocity. The moral ideal of reciprocity can be understood as requiring that all who cooperate in an activity or enterprise must benefit, or share in common burdens, “in some appropriate fashion judged by a suitable benchmark of comparison” (Rawls, 1981:14). The practical requirements of reciprocity differ according to the nature of the collective pursuit under consideration. But some minimal achievement of this ideal is necessary for fairness and effectiveness of collective pursuits of all kinds.

Some reciprocity-based objections to profiling rest on the claim that it fails to distribute the costs of security equally amongst all who benefit but instead imposes them only on groups associated with rule-breaking (Lever, 2004:18). These objections can be met if it can be shown that those who bear the greatest burden are also those who benefit the most. Thus, for example, it is true that the maintenance of a safe school environment is a good enjoyed by society as a whole, because it is conducive to better educational attainment, which in turn produces alumni who have better life chances and are therefore likely to make a more positive contribution (or at least pose less of a threat) to their communities.

But those who benefit most from the maintenance of a safe school environment are the pupils themselves, whose personal life chances are directly affected by decisions such as whether to install metal detectors and introduce spot checks of lockers and bags. And it may well be that the benefits of such measures are worth much more to them than the costs, even when these include stigmatisation and possible reputational damage to the school. This may not always be the case, but when it is the unequal distribution of costs and benefits does not seem obviously unfair (Risse & Zeckhauser, 2004; Lippert-Rasmussen, 2006).⁴

Profiling may be less easy to reconcile with reciprocity when it is used to distinguish between the suspicious and non-suspicious participants in a particular activity or institution. This is partly because stigmatisation costs of all kinds can be particularly severe when people are being singled out for suspicion from among a group of peers. In some cases, these costs may undermine the rules, function or mission of a particular institution or activity more severely than the actions of the rule-breakers they aim to deter. For example, singling out some pupils for surveillance on the basis that they exhibit traits associated

⁴ This is not the same as saying that a society in which some privileged pupils attend good schools and enjoy a safe educational environment while others attend poorly performing schools that subject them to stigmatising and privacy-intrusive surveillance is fair. It is not. But neither is it obvious that such unfairness justifies denying pupils whose education is currently being harmed the measures that might reduce such harm, right here and now.

statistically with higher rates of drug dealing or violence is likely to be deeply humiliating, because teenagers are highly sensitive to the judgements of both their teachers and their peers. Pupils are also highly vulnerable to social exclusion by those peers and to unfair treatment by school authorities, whose marks and reports affect their life chances significantly. Pupils who feel that they are stigmatised by the school authorities or their peers may themselves disengage and withdraw from participation and may be tempted to disrupt the educational environment they feel disfavoured by. Stigmatising some pupils as suspicious may thus have serious consequences for both their willingness and ability to interact well with other students and teachers and to pursue their studies. This may undermine both the social and the educational aims of the school.

Blanket screening of all pupils may be more expensive and disruptive of school routine than profiling, and it may provoke rebellion from pupils who view it as excessively intrusive. But as long as the measure can be shown to be proportionate, the risks to the school's ability to pursue its educational mission effectively are low. Thus there are good reasons to prefer blanket screening to profiling in schools, even when the latter is supported by reliable statistical evidence indicating suspiciousness.

Concerns about whether profiling fulfils the minimal requirements of reciprocity may explain in part why it is often shunned by those responsible for enforcing the rules of particular institutions or activities. For example, all people who apply to work with children in the UK are subject to equal background checks, despite the fact that women are very rarely dangerous to children. This may reflect a desire to ensure that all people who are dangerous to children are deterred from working with them. But it may also reflect a desire to refrain from stigmatising men who want to work with children as potential threats and thus both discouraging qualified and well-intentioned men from applying to work with children. It may also reflect a concern with creating a general mistrust of men and an overconfidence in the intentions of women who want to work with children. This mistrust may affect the nature of professional relationships but it may also have the practical effect of making it harder for men to break the rules of care without being identified but easier for women, thus leading to distortions in the statistics used to justify the use of a profile in the first place (Harcourt, 2004: 1218). It is not difficult to imagine similar kinds of reasoning from reciprocity being used to justify the random drug testing of athletes, the surveillance of all employees, and so on.

2. Untargeted surveillance, the presumption of innocence, and implications of untrustworthiness

Even if these claims about stigmatisation are conceded, it might be argued that there is still a sense in which untargeted surveillance treats people like suspects in ways that conflict with important moral norms. For example, it has been claimed that untargeted surveillance treats innocent people like suspects in ways that undermine the presumption of innocence. It is now generally accepted amongst legal theorists that presumption of innocence is a not

only a legal but also a moral norm whose scope extends beyond the courtroom into security practices more generally (Roberts: 1995).

However, there remain disagreements about the nature and scope of this norm. Traditional accounts of the presumption of innocence understand it as a purely procedural norm: one which protects innocent people from criminalisation by requiring that judgements of innocence or guilt are made by procedures that ensure a sufficient level of certainty. It achieves this primarily by placing the burden of proof on the prosecution and setting the standard of evidence (for most crimes in the EU, beyond reasonable doubt) required to convict. This is the minimum that the presumption of innocence requires on any interpretation of that norm. For traditionalists, this is also the sum of what it requires. Substantive accounts of the presumption of innocence extent view it as requiring that people should be criminalised only if fair procedures have determined them to be guilty of something that can, according to fundamental principles of the law, be legitimately considered a crime (Tadros and Tierny, 2004). For them, the presumption of innocence is violated if fair procedures are used to convict someone of an action that should not have been made illegal in the first place.

The presumption of innocence has also been interpreted as the practical application of the moral principle of civility, which states that all people and institutions “ought to presume, until sufficient evidence is adduced to show otherwise, that any given person has acted in accordance with serious social obligations” (Nance, 1994:648). This interpretation of the presumption of innocence is consistent with both procedural or substantive interpretations. But it extends the presumption of innocence to all areas of social life, including but not limited to the sphere of criminal justice or security practices. Objections to untargeted surveillance from the presumption of innocence reflect all three of the interpretations outlined here.

The first objection concerns the compatibility of untargeted surveillance with the presumption of innocence as a procedural norm. It therefore receives support from all three of the approaches outlined above. This objection states that the necessarily indiscriminate application of untargeted surveillance means it inevitably subjects people who are innocent -sometimes people who are *obviously* innocent - to exactly the same intrusive scrutiny as people whom evidence suggests should be treated like suspects. Because no such evidence exists to support the burdens of scrutiny placed on those innocents surveilled, their imposition is unfair.

One way of responding to this criticism is by drawing a distinction between intruding on people’s privacy in ways that resemble intrusions inflicted on people suspected of wrongdoing, and treating them ‘with suspicion’ or ‘as a suspect’. Not all measures that intrude on privacy or require people to reveal themselves in some way treat people as suspects. Indeed, not all people who come under surveillance by police in criminal investigations are thereby treated with suspicion, although their treatment may resemble in

many ways the treatment of those genuinely suspected of crime or wrongdoing. To illustrate, family members of criminal suspects whose communications are being monitored are often unavoidably also monitored because of their proximity to the suspect. This occurs even when they are not themselves suspected of any crime or collusion in crime. Privacy invasions of this sort are harmful to those they affect and should be avoided. Yet it seems wrong to claim that they treat family members 'with suspicion' or 'as suspects'. It seems wrong even though those individuals are subject to some of the same invasions of privacy as those who are genuinely suspected of a crime. This is because the justification for the imposition of such harms on innocent individuals is not that those individuals are themselves suspicious or fair game for surveillance because they are related to the suspect, but that the imposition of such harms is necessary to the prevention of a greater harm, namely serious crime. The point is that the presumption of innocence does not rule out scrutiny of individuals who are not suspected of any rule-breaking or wrongdoing, even when such invasions appear identical to those inflicted on genuine suspects.

A reasonable condition of the justifiability of such invasions is that they be proportionate, meaning roughly that they are necessary and cause fewer or less serious harms than those of alternatives. In practice, this condition appears to be met in many cases of targeted surveillance. The same can be argued for untargeted surveillance: while the treatment of those surveilled may be similar to that inflicted on suspects, it does not always itself amount to treating people 'with suspicion' or 'as suspects'. It can therefore be justified on grounds other than the existence of evidence that the individuals affected are suspicious. And it can be so without undermining the general rule that evidence of rule-breaking should precede suspicion, not result from it, because not all scrutiny involves suspicion.

This counts against the claim -often made in relation to airport security measures that subject elderly people and children to the same anti-terrorist security checks as nervous-looking young men from Pakistan or the Middle East- that it is both irrational to subject patently innocent or benign people to surveillance (Sir John Stevens in *News of the World*, 2006) and unfair to subject them to the same measures of surveillance as those whom evidence suggests may be a security threat (Lord Brown, 2006, UKHL 12, 77&92).

These points are not contradicted by the existence of real examples of preventive surveillance that *do* threaten to treat people as suspects in ways that undermine the presumption of innocence. The use of counter-terrorism stop and search powers in the UK is one such example. Section 44 of the 2001 Terrorism Act UK allowed police to single people out for body and bag searches on the basis of an undefined suspicion, or 'hunch', that they might be carrying items related to terrorism (Gillian and Quinton vs the United Kingdom ECtHR). The searches were compulsory, meaning that those who refused to submit to them could be subject to imprisonment or fine as a result.⁵ The rationale for this exceptional measure was that the severity of the terrorist threat to the UK provided sound reasons to

⁵ Re Section 44 see (Gillian and Quinton vs UK, ECtHR 33.)

relax the standard of evidence from 'reasonable suspicion' to something less exacting. While this rationale is not objectionable in principle, the resulting policy was: hunches or general intuitions do not qualify as objective grounds for suspicion on any reasonable account of such grounds.⁶ Singling people out for searches on the basis that they exhibit traits that raise suspicion of criminality treats them suspects. Doing so without any objective ground for suspicion undermines the presumption of innocence.

Untargeted surveillance has also been objected to by proponents of a substantive reading of the presumption of innocence. They have argued that it is applied in ways that lead to the criminalisation of individuals on the basis of actions that should not be made criminal in the first place. For example, some argue that the security benefits of untargeted mass databases, such as that proposed under the UK ID card scheme, have not been demonstrated sufficiently (Dept. of Information Systems, LSE, 2005). Yet under some regimes, including that proposed in the UK, failure to register with these schemes results in a fine or other penalty (UK Identity Cards Act 2006: 10(7)). Thus people who do not register and maintain correct details on the database can thereby end up being penalised or criminalised for what in practice amounts to breaking an unjust rule (Chakrabarti, 2004:13; NO2ID in House of Lords, 2009: 27[106]).

It is possible to concede this objection without concluding that untargeted surveillance is any more likely to undermine the presumption of innocence than any other approach to enforcing security. If the UK government have failed to take reasonable measures to demonstrate the probable benefits of ID cards, then the compulsory nature of the scheme may be difficult to reconcile with the substantive presumption of innocence and the more general moral commitment to fair treatment it springs from. But this yields a reason to object to decisions to make participation in schemes of untargeted surveillance compulsory when their benefits are doubtful. It is not a reason to prefer targeted to untargeted surveillance. Neither is it a reason to shun untargeted surveillance as a general approach to enforcing rules and maintaining security.

Sometimes, something like the presumption of innocence appears to be invoked by those who argue that untargeted surveillance involves or is a manifestation of distrust of those surveilled. For example, Norris argues that mass surveillance 'promotes the view...that everybody is untrustworthy. If we are gathering data on people all the time on the basis that they may do something wrong, this is promoting the view that as citizens we cannot be trusted' (Norris in House of Lords, 2009: 27[107]). The criticism deserves close attention, because it articulates an often vaguely expressed concern, which surfaces repeatedly in popular expressions of anxiety about mass surveillance as well as the surveillance studies literature (Lyon, 1994:10; Monahan, 2010:99).

⁶ Lerner (2006) argues that, in practice, the requirement of reasonable suspicion fails to prevent prejudice coming to determine who is subject to a search. Others note that reasonable suspicion is a vague standard vulnerable to questionable interpretation and manipulation by police, and should therefore always be supplemented with clear and detailed guidance (Fundamental Rights Agency, 2010:47)

As mentioned above, some interpret the presumption of innocence as the practical application of the principle of civility. The principle of civility imposes a duty on all to presume, until the contrary has been proven, that people are acting in accordance with their important social obligations. Failure to treat people in ways consistent with this presumption equates to a failure to accord them “dignity associated with the status of membership in the community that is governed by the norms whose breach is at issue” (Nance, 1994: 653). It might be argued that when we treat people as presumptively untrustworthy or suspicious we wrong them, because we cast doubt on their moral status or worth and thus fail to accord them the respect they are owed as equal members of the moral community.

Yet not all presumptive implications of untrustworthiness violate the principle of civility. Some implication of untrustworthiness in individuals is an inevitable consequence of any compulsory mechanisms imposed by authorities to enforce moral and legal rules. In a society whose structure is unjust and which is populated by imperfect individuals, this seems a necessary cost of stability and indeed of justice. This is not to deny that the principle of civility should inform our interactions with each other, including interactions involving surveillance. It is merely to point out that the principle of civility is not the only or the most important norm governing the use of surveillance in the enforcement of moral rules.

Sometimes some implication of mistrust is necessary if people are to be able to be and feel secure enough to pursue practices important or even crucial to their wellbeing. Sometimes the risks and fear of harm associated with an activity are high but the evidence enabling authorities to prevent that harm in a targeted way is lacking. In such cases, some measure of untargeted surveillance may be a necessary condition of people being and feeling secure enough to continue to engage in the activity. For example, it is likely that a number of people not feel secure enough to continue driving cars or travelling by plane if the blanket use of speed cameras and of passenger and luggage screening were to cease.

It seems reasonable to conclude that the implication of mistrust involved in surveillance need not violate the principle of civility *unless* individuals are singled out for suspicion on the basis of insufficient evidence and unless they are surveilled more than is necessary and proportionate to the protection of their (democratically or otherwise legitimately defined) collective interests. Arguments against unnecessary and disproportionate impositions of surveillance are arguments in favour of fair and effective security policies. They are not arguments against untargeted surveillance *per se*.

None of what has been argued here suggests that untargeted surveillance never conflicts with or undermines the presumption of innocence. But it does suggest that untargeted surveillance need not conflict with the presumption of innocence any more than targeted surveillance.

3. Untargeted surveillance and privacy

While critics of untargeted surveillance overstate its suspicion-based risks, they overlook its privacy risks, broadly understood. In this section I explain how untargeted surveillance can affect people's privacy.

Privacy is an important value for individuals and, as is increasingly recognised by ethical theorists, it also has an important social value. Individuals need privacy to build and maintain meaningful relationships, to express their feelings and desires freely in artistic and political ways, and to experiment with and arrive at ideas for themselves about how they want to live their lives. Thus privacy is a condition of both personal happiness and individual freedom (Sorell, 2011). Privacy is also a condition of a functioning liberal democracy. Without a private space in which to express and exchange political ideas and opinions, explore and practice religious beliefs, teach one's children one's own values and vote anonymously, amongst other things, people's ability to engage in activities of democratic citizenship with genuine autonomy, that is, free of exploitation or oppression, would be weakened (Lever, 2011; Solove, 2008; Allen 2011). This, in turn, would weaken the effectiveness of democracy for society as a whole. For these reasons at least, privacy should be treated as an important value or freedom and should be limited or interfered with only to the extent that is proportionate to the protection of other equally or more important values or interests.

There is some disagreement amongst scholars as to which kinds of surveillance interfere with people's privacy and which affect other interests. For example, it has been argued that CCTV installed in public areas neither invades any recognised zone or sphere of privacy, nor (at least typically) gathers information of a personal or sensitive nature (Ryberg, 2007). On the other hand, it has been argued that information about a person's location can itself count as personal (and therefore presumptively private) information (van der Hilst, 2010). It has also been argued that the question of whether a certain measure affects privacy interests cannot be answered without consideration of the identity of the agent doing the watching (Lever, 2008). Most measures considered in this paper, and indeed most kinds of surveillance whose aim is to prevent rule-breaking behaviour in general, do affect people's privacy interests, at least on most understandings of those interests. And even those forms of surveillance which, like some forms of CCTV, scrutinise people in public places, can be said to interfere with privacy when such scrutiny chills the pursuit of permitted private activities such as intimate conversations, physical contact, political debate and so on.

Interferences with privacy weigh more heavily in considerations of the risks of surveillance the more fundamental to freedom or well-being the interest they affect, the more intrusive they are and the greater the number of people they affect. Intrusiveness is itself determined by a number of factors. These include the extent to which the interference is reasonably to be expected, can be planned for, and is consented to; the number of people given access to

the private information or zone; the sensitivity of the information or zone affected; and the period for which any data gathered is retained.

Unlike targeted surveillance, which typically intrudes quite seriously on the privacy of few, selected targets, untargeted surveillance is often unavoidably intrusive to large numbers of individuals. In addition, untargeted surveillance almost inevitably intrudes on far greater numbers of innocent individuals than more targeted alternatives. The contrast is most stark between untargeted surveillance that affects entire populations on the one hand, and those targeting an individual suspect; profiling and random screening fall somewhere in between.⁷

The fact that untargeted surveillance often infringes the privacy of large numbers of people is an inevitable outcome of the indiscriminate nature of the surveillance. However, the intrusion visited on each individual is in practice often mitigated by the fact that the surveillance tends to be overt and well-publicised, allowing people to adapt their expectations and plan accordingly. This is the case with airport screening, speed cameras, and the surveillance of individuals working with sensitive chemicals or children, for example. Some intrusions, such as the in-depth investigations of high-ranking government officials, or the drug testing of athletes, are not minor. But most are consented to, albeit sometimes only tacitly. Exceptions include measures such as the collection of biometric information for criminal justice databases, which in many cases cannot be avoided or cannot be avoided without incurring legal penalty.

4. Surveillance and discrimination

A further risk associated with surveillance that should be taken into account in any assessment of its potential costs is discrimination. Discrimination in surveillance can occur when people are singled out for scrutiny on the basis of factors other than what makes them suspect. Much more commonly, however, it occurs when surveilling authorities deciding whom to single out place too much weight on factors only weakly or indirectly related to people's suspiciousness. For example, airport security that singled out everyone travelling to Europe from Pakistan for extra security checks on the basis that people travelling from Pakistan are more likely to have attended terrorist training camps is probably discriminatory. Millions of people travel from Pakistan to Europe every year. People who travel to Europe from Pakistan might be more likely statistically to have attended a terrorist training camp, but this correlation is far too weak to justify singling them all out for extra security checks. The likelihood is small that any such checks will deter or detect terrorists effectively enough to begin to justify the extra burden on airport security and the stigmatisation and privacy costs that will inevitably result from singling them out.

Discrimination is less likely to occur with the most and least targeted forms of surveillance than it is with those falling somewhere in between. This is because both entirely untargeted

⁷ For a similar argument about the relative privacy costs of targeted and 'dragnet' searches in the USA see Primus 2011:6.

and highly targeted surveillance afford surveilling authorities very limited discretion to make generalisations about the kinds of people likely to be involved in unwanted behaviour. For example, airport security measures that affect all passengers equally do not rely on generalisations about which kinds of passengers are more suspicious and therefore do not run a high risk of discriminating against one particular group. At the same time, measures that target only individuals whom police reasonably suspect of criminality are also less likely to be discriminatory, given the considerable strength and objectivity of the evidence needed to meet the threshold of reasonable suspicion. In contrast, profiling by police is more likely to be discriminatory, given the fact that it often relies on broad generalisations about which traits are indicative of suspiciousness.

The risk of discrimination also rises the more discretion is given to individual surveillance officers, because all individuals harbour witting and unwitting prejudices, which, given the opportunity, they may act on. For example, discrimination against Asian and Muslim people appears to have resulted from a programme of anti-terrorism stop-and-search in the UK under Section 47A (Moeckli, 2007; Bowling and Philips, 2007; Parmar 2010) of the 2001 Terrorism Act. As mentioned above, this Act empowered police to stop individuals on the basis of 'hunches' or 'professional instinct'. Perhaps unsurprisingly, it appears that in practice the hunches of UK metropolitan police often reflect popular prejudices about the criminal propensities of certain ethnic groups. It seems reasonable to expect that, had the standard of reasonable suspicion been maintained, discrimination would have occurred to a lesser extent, because police would have been obliged to present objective evidence demonstrating suspiciousness. Had police been instructed to stop everyone passing a certain point, in something like a blanket search, the ability of police to act on prejudices would have similarly been limited, thus also reducing the risk of discrimination.

5. Surveillance and mission creep

Before moving on to consider what the relative benefits of untargeted and targeted surveillance are and whether and when these might justify the stigmatisation and privacy intrusions incurred, it is important to factor in the risk of mission creep. Mission creep occurs when data collected through surveillance is used for a purpose other than that for which it was originally approved. Surveillance based on individual suspicion raises the lowest risk of mission creep. Profiling and untargeted surveillance raise similar risks of mission creep, because they more commonly involve the storage of large amounts of data for future use. Concerns about mission creep in untargeted surveillance have been raised in relation to the use in the UK of ANPR data to identify individuals attending or presenting in the vicinity of protests (Guardian Newspaper, 2012), and in relation to the collection of biometric material in medical and other databases. Privacy, stigmatisation, and discrimination risks not typically associated with untargeted surveillance may nevertheless result from it if it is used as a precursor to target individuals for suspicion. Mission creep is unjustified when the new purpose for which the data is used has not been approved

through a legitimate democratic process, or when the oversight provided by such processes is poor.

6. The relative moral risks of targeted and untargeted surveillance reassessed

Taken together, the claims made here about the relative risks of untargeted and targeted surveillance support the following position: the more targeted a measure of surveillance, the more likely it is to stigmatise those it singles out for scrutiny. The less targeted a measure of surveillance, the more likely it is to impose privacy costs on large numbers of people and on large numbers of obviously benign or innocent individuals.

To illustrate these claims, a graph has been inserted in fig 1 below. This graph is intended to be a rough indication of the relationship between the most serious risks of surveillance used to reduce unwanted behaviour and the extent to which that surveillance is targeted on the basis of suspiciousness.

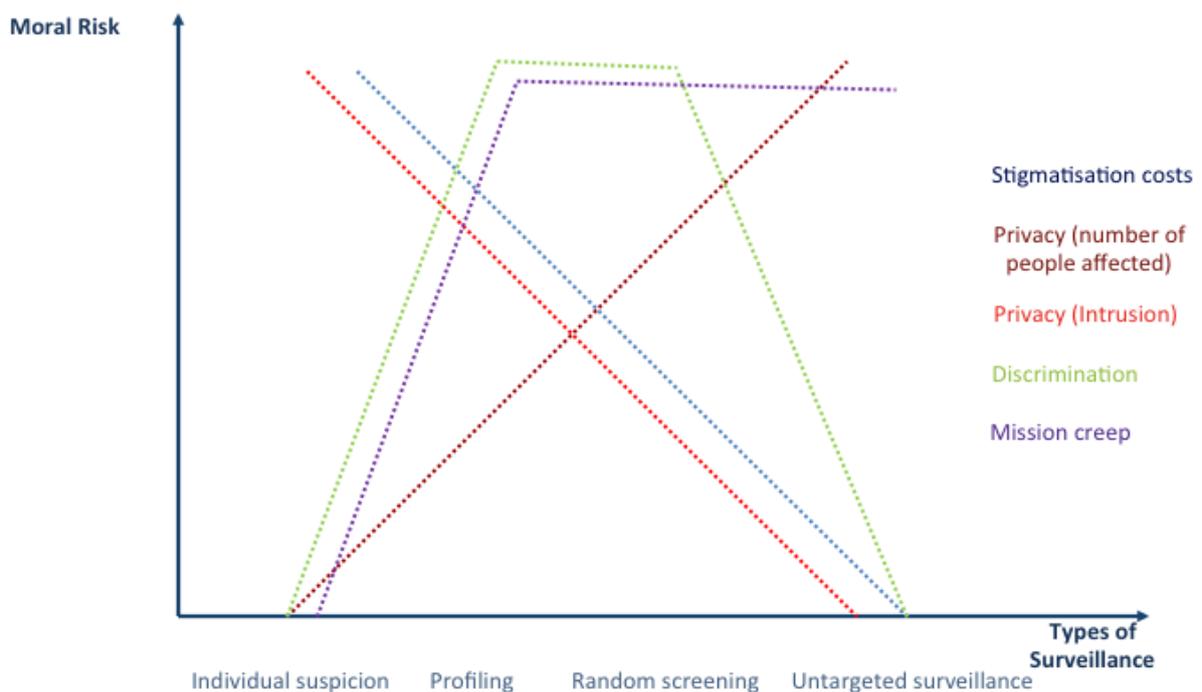


Figure 1

Measures that fall somewhere between targeted and untargeted surveillance often pose some of the risks or incur some of the costs of both those extreme measures. For example, random screening singles people out and so carries greater risks of stigmatisation than blanket screening. It also carries greater risks of discrimination, as prejudice may come to influence how decisions about whom to suspect are made. However, random screening also imposes privacy costs on fewer individuals than blanket screening. It also poses a smaller risk of stigmatising costs than actions that single people out on the basis of individual suspicion (as long as there is public awareness of the random basis for the selection). On the other hand, measures like profiling that are targeted at suspicious categories carry high risks

of both stigmatisation and discrimination associated with targeted surveillance. They also inevitably result in a high proportion of false positives. These are some of the reasons why measures like profiling have proved highly controversial in practice.

Some security practices combine targeted and untargeted surveillance in ways that mean their costs are more varied than those of straightforwardly targeted or untargeted alternatives. Readers who are very familiar with the surveillance studies literature may well have noted the conspicuous absence of open street CCTV from my lists of examples of untargeted surveillance. This is because open street CCTV is often used in both a targeted and untargeted way and, up until now, my intention in this paper has been to highlight rather than to play down the distinctions between these two approaches. There are two stages to the use of open street CCTV, at least as it is currently used in the UK. The first is untargeted: open street CCTV affects all people who pass through its range by recording them.⁸ The purpose of this stage of CCTV is to keep a record of events that can be accessed if an incident of rule-breaking occurs. The second stage of CCTV use is highly targeted: once an incident occurs, state agents watch only those recordings that independent information suggests may hold evidence that will help them to establish the facts of a case.

Open street CCTV incurs risks at both ends of the surveillance spectrum. The privacy, stigmatisation and discrimination risks of open street CCTV are borne primarily by those who are targeted for scrutiny in relation to an incident of rule-breaking. Having one's movements scrutinised by state agents is far more intrusive than having them recorded but then deleted, unwatched by anybody. No one is stigmatised or discriminated against by being recorded, because no one is singled out. Neither is the interference presented by the recording itself significant. Indeed the costs of the untargeted use of CCTV consist primarily in the risk to individuals that their public actions will be scrutinised if the law is broken in a location in which they also happened to be. The costs to each individual of untargeted CCTV are therefore very low, but they are costs that must be borne by vast numbers of individuals.

Whenever proposals to establish a new programme of surveillance are made, the moral risks arising from both targeted and untargeted aspects of its use should be taken into account. Open street CCTV is not unique in combining aspects of both. Large-scale databases may similarly incur moral risks at both ends of the spectrum: having one's details stored on a database is far less invasive and stigmatising than having them scrutinised and connected up with other details stored on other databases, although the former is often only undertaken in order to enable the latter.

The utility of the scale in Fig.1 above lies in its ability to help us to understand what the main moral risks of surveillance are, when they are most likely to arise, and how trade-offs between them might be made. But there remain morally significant aspects of surveillance

policy that it is too simplistic to reveal. For example, some people with prostheses, disabilities or other special circumstances may be forced by metal detectors and body scanners of some kinds to reveal these intimate aspects of themselves, aspects which they might reasonably desire to conceal. Thus the scrutiny of untargeted surveillance may be indiscriminate, but the intrusion into privacy may be greater for some individuals than others.

6. The benefits of surveillance: distinguishing between deterrence and prosecution

Just as the potential costs of targeted and untargeted forms of surveillance differ, so do the benefits. This is often overlooked by critics of untargeted surveillance, who tend to assume that the only measure of success of a security policy is the number of threats detected or wrongdoers apprehended. In this section I argue that this wrongly excludes deterrence, which is both a legitimate and valuable aim and most often the principal objective of untargeted surveillance.

It has been argued that untargeted surveillance is less effective in producing benefits to security than more targeted surveillance, including profiling, because it results in far fewer arrests or convictions (Risse and Zeckhauser, 2004, REF). But this ignores the fact that security is protected by deterring people from breaking the rules, as well as by identifying those who have already broken them. It is only the effectiveness of a practice in identifying rule-breakers that can be measured by reference to the number of arrests made. In contrast, its effectiveness in deterring crime can be measured by comparing the number of incidents of rule-breaking reported prior to and following its implementation. If the number has fallen then, other things being equal, the measure has had some positive effect.

Most untargeted surveillance used to enforce the rules of specific institutions or particular activities aims at deterring rule breakers rather than identifying them. Thus an absence of positive identifications of rule-breakers is not a sign of failure. This is true even of measures whose value as a deterrent is their reliability in catching people out. For example, metal detectors at airports are intended to detect any and all metal items carried on an individual's person and thereby to deter individuals from carrying weapons or other dangerous items onto airplanes. If people continued to carry weapons onto planes despite the use of metal detectors, this would signal that metal detectors were failing in their aim of deterring people from hazardous behaviour. Similarly, the expectation when undertaking surveillance of employees or of people applying to work with children or of athletes and so on is that knowledge of the imminent scrutiny will discourage those who intend to break the rules from participating in the activity and thereby prevent their violation. There are exceptions, including open street CCTV and the use of large databases, both of which often enable more targeted surveillance whose aim is to identify rule-breakers. In these cases, the rate of positive identifications is an appropriate measure of effectiveness. But for most

forms of untargeted surveillance, it seems unreasonable to insist that the low rate of positive identification or detection means its security benefits are non-existent.⁹

Other things being equal, measures of untargeted surveillance are more effective at both revealing and deterring rule-breaking than profiling or measures that rely on individual suspicion, because they are very difficult to evade. However, in practice it is likely to be most efficient when it is used to enforce the rules of a well-defined, circumscribed, and easy-to-monitor activity or institution. Surveillance that aims to maintain the rules of specific institutions or activities is much less costly to carry out in an untargeted way than surveillance that aims to enforce the rules of the criminal law, which can be broken anywhere, at any time, by anyone and in a vast array of ways.¹⁰ In order to be effective, untargeted surveillance that attempted to deter violations of the criminal law in general would have to be as pervasive as the network of surveillance described in George Orwell's 1984, which intruded in all areas of life including those considered most private. The costs in terms of money and manpower of installing and maintaining such a system would probably be far greater than those involved in running the current, more reactive approach to criminal justice. Thus the most untargeted forms of surveillance are likely to be an inefficient means of enforcing the criminal law.

Just as completely untargeted surveillance is likely to be inefficient at enforcing criminal justice, so the most targeted forms of surveillance, i.e. those that require individual suspicion to trigger scrutiny, are often inefficient at enforcing the rules of specific institutions or activities. In many cases there is a lack of available evidence about who is likely to break the rules, and who can be ruled out from suspicion. In many cases the costs of gathering information and following up evidence about specific people in advance of any incident of rule-breaking may be very high. Yet both doing nothing and relying on profiles, which are never impossible to evade, may expose vulnerable people to unacceptable security threats or valuable institutions to being undermined through rule-breaking or wrongdoing. Searching everyone who comes through an airport may be more efficient at maintaining security than trying to identify which of those travellers are likely to be threats and targeting only them.

These arguments can be usefully illustrated by thinking about airport security. The maintenance of law and order on aeroplanes is potentially more difficult than it is in most other environments. Planes are cramped places in which people are forced to share small amounts of space. For this reason, disputes that turn violent are impossible for other

⁹ Bernard Harcourt has argued convincingly that the detection or 'hit rate' is also an inaccurate measure of the effectiveness of profiling, because it fails to track both the 'overall amount of profiled crime and the costs to society of searches'. (Harcourt, 2004:1281).

¹⁰ Some of the rules enforced under the auspices of specific institutions overlap with rules of criminal justice. Thus dealing drugs breaks both school rules and the law. However, the primary aim of spot-checks of bags and lockers in schools is to maintain the school rules and thereby ensure a sound education environment, rather than to enforce the law.

passengers to escape or distance themselves from. Because planes take time to land, any injury caused as a result of disorder caused while airborne could only be treated with significant delay. Any disturbance involving the stewards could compromise the security of the rest of the passengers and if this spread to the cockpit the consequences could be much worse. For all of these reasons, it is permissible for security agencies to take some intrusive measures to deter or to chill hazardous behaviour.

Metal detectors and hand-luggage x-rays are minimally intrusive. First, they are used in airports, which are controlled environments in which reasonable expectations to privacy are diminished. For example, people using airports are given ample warning of the fact that they are likely to be screened. This enables people to alter their behaviour accordingly, to ensure, for example, that items they do not want to be seen by others are excluded from their hand-luggage. Second, the searches are not compulsory because travel by plane is not compulsory. Though for some people it is a necessary aspect of their job or family life, for most people traveling by plane is something they can forego without hindering the free pursuit of their interests.¹¹ Third, travellers are usually willing to consent to such searches because they receive important benefits in terms of both security and reassurance that their security is protected. Given that many people feel afraid of the security risks involved in flying, the latter helps, in turn, to maintain order while in flight.

If security agencies were required to act only on the basis of evidence about individuals or the kinds of people likely to be threats to airport security, the intrusions visited on individuals travelling by plane would be likely to be far greater. It is difficult to predict who is likely to be carrying a weapon, for example, without gathering detailed information about them. The process of gathering such information is likely to be more intrusive than asking them to step through a metal detector. It is also likely to be stigmatising, as travellers are subject to different kinds of intrusion depending on their suspiciousness. Moreover, if people knew that they would only be searched at airports if police had some evidence or intelligence suggesting wrongdoing, they would be more likely to engage in hazardous behaviour such as carrying knives or guns on their person, even if they did not intend to use them. Thus we can begin to see why untargeted surveillance is used when it is and why in those cases it may be a more efficient and less morally costly means of achieving important security aims than targeted alternatives.

Conclusion

The aim of this paper has been neither to put forward an unqualified defence of untargeted surveillance nor to suggest that it is always or in most cases less risky morally than targeted alternatives. Rather, the aim has been to point out some shortcomings in current

¹¹ Marx points out that sometimes the choice to opt out from untargeted surveillance comes at a high cost to individuals. He claims that in these cases the extent to which the choice can be said to be voluntary is questionable (Marx, 2007:16). This seems correct, but while genuinely voluntary consent might make surveillance easier to justify it is not a necessary condition of its being justifiable.

assessments of the moral risks and the benefits of targeted and untargeted approaches to surveillance, and to propose a more systematic approach to understanding both.

While some general claims can be made about the tendency of different kinds of surveillance to generate stigmatisation, privacy, and discrimination costs when used for certain kinds of purposes, very little can be deduced from what has been said here about the justice of any individual measure. For that, much more information is needed about the specific circumstances in which the surveillance will be implemented, not to mention an account of what justice consists in. Until now, most moral theorists have adopted either contractarian or consequentialist approaches to analysing security policy. The scope of this paper does not extend to a defence of one such theory over any other, but the proposed analysis of the costs of surveillance aims to be of equal relevance to both.¹²

List of References

Allen Anita, *Unpopular Privacy*, 2011

Arneson Richard, 'Shame, Stigma and Disgust in the Decent Society' in *Journal of Ethics* 2007

BBC, 'Concern over 2012 athletes' village drug tests', Jan 4, 2010 at http://news.bbc.co.uk/sport1/hi/olympic_games/london_2012/8440088.stm

Bou-Habib Paul, 'Profiling, Security and Equality' in *Ethical Theory and Moral Practice* 2008

Bou-Habib Paul, 'Racial Profiling and Background Injustice' in *Journal of Ethics*, 15(1-2), 2011

Bowling, Ben and Phillips, Coretta, 'Disproportionate and Discriminatory: Reviewing the Evidence on Police Stop and Search' in *The Modern Law Review* Vol.70.6 (2007)

Clancy, Thomas, 'The Role of Individualized Suspicion in Assessing the Reasonableness of Searches and Seizures' in *The University of Memphis Law Review*, Winter. 25. 1994

Courtwright 'Stigmatisation and Public Health Ethics' in *Bioethics*, 2011

¹² The author would like to thank Tom Sorell, John Guelke, Kevin Macnish, Rose van der Hilst and Mathias Vermeulen for their comments and advice.

Department of Information systems, LSE. *The identity project : an assessment of the UK Identity Cards Bill and its implications*. Department of Information Systems, London School of Economics and Political Science, London, UK. (2005) At <http://eprints.lse.ac.uk/684/>

Fundamental Rights Agency of the EU (FRA), *Towards More Effective Policing. Understanding and Preventing Discriminatory Ethnic Profiling: A Guide*, 2010 at http://fra.europa.eu/fraWebsite/attachments/Guide_ethnic_profiling.pdf

Gillian vs Quinton, European Court of Human Rights, Judgement, Strasbourg 12 January 2010

Grayling A. *Liberty in the Age of Terror* (London: Bloomsbury, 2009)

Guardian Newspaper 'Protester sues police over surveillance database', Feb 9th 2012. At <http://www.guardian.co.uk/uk/2012/feb/09/protester-sues-police-surveillance-database>

Haggerty and Ericson, *Policing the Risk Society*, 1997

Harcourt Bernard, 'Rethinking Racial Profiling: A Critique of the Economics, Civil Liberties, and Constitutional Literature, and of Criminal Profiling More Generally' *University of Chicago Law Review*, 71(4) 2004

Human Rights Watch *Without Suspicion: Stop and Search under the Terrorism Act 2000*, 2010

House of Lords, Constitution Committee. *Surveillance and the State: Second Report, 2008-9*. Jan 2009. <http://www.publications.parliament.uk/pa/ld200809/ldselect/ldconst/18/1802.htm>. Last accessed 1st July 2012.

Kennedy Randall, "Race, Law and Suspicion" in *Race, Crime and the Law*, (New York: Vintage Books, 1997)

Lerner, Craig S. 'Reasonable Suspicion and Mere Hunches' *Vanderbilt Law Review* 59:2, 2006

Lever Annabelle, 'Why Racial Profiling is Unjustified' *Philosophy and Public Affairs* Vol. 32, No. 2, Spring 2004

Lever, Annabelle, 'Mrs Aremac and the Camera: A Response to Ryberg' *Res Publica* 14(1) 2008

Lever Annabelle, *Privacy: a Short Introduction*, 2011

Liberty, *From War to Law: Liberty's response to the coalition government's review of counter-terrorism legislation*, 2010

Lippert-Rasmussen, 'Racial Profiling versus Community', *Journal of Applied Philosophy*, Vol. 23, No. 2, 2006

Lyon, David, *The Electronic Eye: The Rise of Surveillance Society*. University of Minnesota Press, 1994

Marx, G, 'Hey Buddy, Can you Spare a DNA? New Surveillance Technologies and the Growth of Mandatory Volunteerism in Collecting Information' in *Ann Ist Super Sanita* Vol.43(1), 2007.

Moeckli Daniel, 'Stop and Search Under the Terrorism Act 2000: A Comment on R (Gillan) v UK' in *Modern Law Review* 2007, 70(4)

Monahan, Torin 'Surveillance as Governance: Social Inequality and the Pursuit of Democratic Surveillance' in Haggerty and Samatas (eds), *Surveillance and Democracy*, Routledge, 2010

Nance, Dale A, 'Civility and the Burden of Proof' *Harvard Journal of Law and Public Policy* 17 (1994)

New Scientist. 'All Seeing Eyes' Thursday, November 02, 2006 at <http://www.newscientist.com/blog/technology/2006/11/all-seeing-eyes.html>. Last Accessed 1st July 2012.

O'Connor Michael and Rumann Celia, 'Into the fire: How to avoid getting burned by the same mistakes made fighting terrorism in Northern Ireland', 24 *Cardozo L. Rev.* 1657 (2003).

Parmar, Alpa 'Stop and Search in London: Counter-Terrorist or Counter-Productive?' *Policing and Society* 21:4, 2012

Primus, 'Disentangling Administrative Searches' in *Columbia Law Review*, 2011

Rawls, 'The Basic Liberties and their Priority', *The Tanner Lectures*, University of Michigan, 1981

Reiman, 'Is Racial Profiling Just? Making Criminal Justice Policy in the Original Position' in *Journal of Ethics*, 15(1-2), 2011

Report of the UK Independent Reviewer of Counter-Terrorism 2010 Lord Carlile March 2010.

Risse and Zeckhauser, 'Racial Profiling' in *Philosophy and Public Affairs* Vol. 322, No. 2, (Spring 2004)

Roberts Paul, 'Taking the Burden of Proof Seriously' *Criminal Law Review* 1995

Ryberg, 'Privacy Rights, Crime Prevention, CCTV, and the Life of Mrs Aremac' in *Res Publica* 2007

Schauer, 'Generality and Equality', *Law and Philosophy* Vol.16, 1997

Solove, *Understanding Privacy*, 2008

Sorell, 'Preventive Policing, Surveillance, and European Counter-Terrorism' in *Criminal Justice Ethics*, 2010

Sir John Stevens 'If you're Muslim, it's your Problem' in *News of the World*, August 13, 2006

Tadros Victor and Tierny Stephen, 'The Presumption of Innocence and the Human Rights Act' in *Modern Law Review*, 2004 67(3).

UK House of Lords 12, 2006: Judgments - R (on the application of Gillan (FC) and another (FC)) (Appellants) v. Commissioner of Police for the Metropolis and another (Respondents)

Van der Hilst, Rozemarijn, 'Characteristics and Uses of Selected Detection Technologies: Tracking Technologies'. DETECTER Project Deliverable 17.3. 2011. At www.detecter.eu

The Relative Moral Risks of Untargeted and Targeted Surveillance

Abstract: Is surveillance that is targeted towards specific individuals easier to justify than surveillance that targets broad categories of people? Untargeted surveillance is routinely accused of treating innocent people as suspects in ways that are unfair and of failing to pursue security effectively. I argue that in a wide range of cases untargeted surveillance treats people less like suspects than more targeted alternatives. I also argue that it often deters unwanted behaviour more effectively than targeted alternatives, including profiling. In practice, untargeted surveillance is likely to be least costly morally and most efficient as a means of enforcing the rules of a specific activity or institution. Targeted alternatives are likely to be more appropriate means of enforcing the rules of criminal law in general.

Keywords: Surveillance; privacy; stigmatisation; discrimination; reciprocity