

Risks, Rights, Statistics and Compulsory Measures

DENISE MEYERSON*

Abstract

Preventive coercive measures are an increasingly common feature of the legal landscape. Such measures are based on the probability, not the certainty, of future harmful conduct and they therefore carry the danger that some people who are not dangerous will be mistakenly subjected to them. Such people are called ‘false positives’. This article explores the moral and legal complexities which surround the problem of false positives. It investigates two lines of reasoning which attempt to minimise the severity of the problem. Explaining why we should reject these approaches, the article argues that, as a matter of political morality, it is more important to avoid mistaken decisions to restrict liberty than mistaken decisions not to restrict liberty. It also explains the nature of the legal protections which are needed to give effect to this weighting of the competing values at stake. Finally, it sheds light on provisions prohibiting arbitrary detention and arbitrary invasions of privacy in recent human rights legislation, explaining how the view adopted in this article can be used to give meaning to the notion of ‘arbitrariness’.

1. Introduction

The traditional backwards-looking, reactive response of the law to harmful conduct — the imposition of liability on the basis of past conduct — is shifting and emphasis is increasingly placed on preventing harmful conduct before it occurs. This is part of a more general movement to what has been aptly called a ‘risk society’, a society increasingly preoccupied with safety and the future.¹ On this safety-dominated, forwards-looking approach, the law may restrict the rights and liberties of individuals not because of what they have done in the past, and not even for a mixture of preventive and non-preventive reasons, but purely to protect society against the perceived risk of certain individuals causing harm in the future.

One site for these developments is the criminal justice system. Criminologists have observed a shift from the traditional approach to criminal justice, which is structured around the ideas of morality, guilt and individual responsibility, to a new, actuarial discourse, which employs the language of probability and statistics

* Professor, Macquarie Law School, Macquarie University. I would like to thank Katherine Biber for helpful conversations on these matters.

1 Ulrich Beck, *World Risk Society* (1999); Anthony Giddens, ‘Risk and Responsibility’ (1999) 62 *Modern Law Review* 1, 3.

and aims to manage, control and incapacitate members of high-risk groups.² We see this trend at work in Australian sentencing laws which permit longer than proportionate sentences or indefinite sentences in order to protect society against offenders who are thought likely to commit further crimes. The same rationale informs laws which license the preventive imprisonment of dangerous individuals after the expiration of their sentences. In both cases, the laws aim purely at incapacitation of the dangerous person. The former do so in respect of that portion of the sentence which cannot be justified by desert. The latter do so in respect of the whole period of imprisonment after expiration of the deserved sentence, since no new crime has been committed.³

Other Australian laws provide for intrusive restrictions on the liberty and privacy of sex offenders who have completed their sentences and have been released into the community but are thought likely to commit further offences. These may include long-term limitations on their freedom of movement, place of work and residence, and subjection of the offender to electronic monitoring, curfews and restrictions on activities such as use of the internet.⁴ In some Australian jurisdictions, outlaw motorcycle gangs have also recently become the target of legislation aimed at reducing the threat of serious and organised crime. This legislation allows for the making of control orders against gang members so as to restrict their activities and prevent them from associating with other gang members, and for the issuing of public safety orders which may prohibit gang members from, for instance, attending a public event or place.⁵

The mental health context provides further examples. Although involuntary commitment to a mental hospital is generally justified at least in part on the ‘best interests’ or *parens patriae* ground that the person will benefit from treatment, some Australian mental health laws appear to permit involuntary commitment purely on the public safety ground that control of the person is necessary in order to prevent serious harm likely to flow from their mental illness or mental disorder.⁶

2 Malcolm Feeley and Jonathan Simon, ‘Actuarial Justice: The emerging new criminal law’ in David Nelken (ed), *The Futures of Criminology* (1994) 173; Malcolm Feeley and Jonathan Simon, ‘The New Penology’ in John Muncie, Eugene McLaughlin and Mary Langan (eds), *Criminological Perspectives: A Reader* (1996) 367.

3 For discussion of the relevant legislation, see Ben Power, ‘“For the Term of his Natural Life”: Indefinite sentences — a review of current law and a proposal for reform’ (2007) 18 *Criminal Law Forum* 59.

4 For discussion of Victoria’s sex offender monitoring legislation — the *Serious Sex Offenders Monitoring Act 2005* (Vic) — see Mark Brown, ‘Risk, Punishment and Liberty’ in Thalia Anthony and Chris Cunneen (eds), *The Critical Criminology Companion* (2008) 253, 257–9.

5 *Serious and Organised Crime (Control) Act 2008* (SA) pts 3 and 4; *Crimes (Criminal Organisations Control) Act 2009* (NSW) pt 3.

6 Sections 14 and 15 of the *Mental Health Act 2007* (NSW), for instance, use the language of ‘care, treatment or control’. Terry Carney, David Tait and Fleur Beaupert comment that ‘[t]his formulation suggests that need for control alone to avert dangerousness without a baseline element of need for treatment satisfies these prerequisites [for civil commitment]’: ‘Pushing the Boundaries: Realising rights through mental health tribunal processes?’ (2008) 30 *Sydney Law Review* 329, 339. See also Bernadette McSherry, ‘“Dangerousness” and Public Health: Civil detention of individuals with infectious diseases’ (1998) 23 *Alternative Law Journal* 276, 278.

Furthermore, there are increasing calls on the mental health system to protect society from those who suffer from personality disorders such as antisocial personality disorder or ‘psychopathy’. Although people with antisocial personality disorder can be extremely dangerous, it is generally accepted among psychiatrists that they are neither mentally ill nor suffering from a mental disorder and that they are very unlikely to benefit from treatment.⁷ There is resistance in Australia to using the mental health system to incapacitate such people,⁸ but in the United States the idea has been openly embraced in the form of the so-called ‘Sexually Violent Predator laws’. These laws authorise the post-imprisonment indefinite civil commitment of sex offenders who have a ‘mental abnormality’ or ‘personality disorder’ which makes them likely to engage in repeat acts of sexual violence. A mental disorder or illness is not required.⁹

It is perhaps worth explaining that ‘mental abnormality’ is not a concept recognised within the discipline of psychiatry,¹⁰ and it appears that the Sexually Violent Predator statutes make reference to it merely in order to ward off constitutional difficulties.¹¹ They formally provide for treatment for the same reason but this is not their goal and it is recognised that the treatment is not likely to be and does not have to be efficacious.¹² Such laws have led some to express the fear that psychiatry will ‘assume an Orwellian air, as the socially undesirable risk indefinite incarceration in psychiatric (or pseudo-psychiatric) institutions’¹³ and doctors are ‘required to pretend to treat the untreatable for the sake of a third party.’¹⁴

In the anti-terrorism context, there are Australian laws which seek to protect the public from terrorist acts. These laws authorise the detention of individuals without charges and the imposition of control orders on individuals who have not been charged with terrorist crimes. Such control orders can, like the orders placed on released sex offenders and gang members, impose severe restrictions and surveillance on the individuals subject to them.¹⁵

7 C R Williams, ‘Psychopathy, Mental Illness and Preventive Detention: Issues arising from the David case’ (1990) 16 *Monash University Law Review* 161, 165–6.

8 C R Williams discusses the failed attempt in Victoria to use the mental health system to preventively detain Garry David after the expiry of his sentence: *ibid.*

9 Such laws were upheld by the United States Supreme Court in *Kansas v Hendricks*, 521 US 346 (1997) and *Kansas v Crane*, 534 US 407 (2002). For discussion of the laws, see Roxanne Lieb, Vernon Quinsey and Lucy Berliner, ‘Sexual Predators and Social Policy’ (1998) 23 *Crime and Justice* 43, 65–9.

10 Alexander D Brooks, ‘The Constitutionality and Morality of Civilly Committing Violent Sexual Predators’ (1992) 15 *University of Puget Sound Law Review* 709, 730.

11 In *Foucha v Louisiana*, 504 US 71 (1992) the United States Supreme Court struck down a statute that permitted the indefinite confinement in a mental institution of dangerous defendants who had been acquitted by reason of insanity but were no longer mentally ill.

12 Brooks, above n 10, concludes at 752: ‘the bottom line is that, whether treatment works or not, the US Supreme Court has constitutionally validated confinement that is designed to protect society against mentally abnormal dangerous persons.’

13 Frank R Farnham and David V James, ‘“Dangerousness” and Dangerous Law’ (2001) 358 *The Lancet* 1926.

14 *Ibid.*

15 For details of the legislation, see Andrew Lynch and George Williams, *What Price Security?: Taking stock of Australia’s anti-terror laws* (2006) 41–58.

Finally, all Australian jurisdictions have laws which permit the detention of contagious individuals and the forcible quarantining of healthy people who have been exposed to an infectious disease.¹⁶ The potential problems with these laws tend to be largely overlooked because they are so little used. They have generated controversy, however, in relation to the detention of people with AIDS,¹⁷ and the difficulties may start to loom large if the fear of bio-terrorist attacks materialises.

2. *The Problem of False Positives*

The laws referred to above are offered as examples only. I have made no attempt to be comprehensive, or to provide details of the relevant legislation, because my interest is not in documenting the law. It is in the moral problems attached to serious curtailments of liberty for no reason other than the good of others. It is obvious that such measures pose numerous problems of political morality, not least by eliciting our great fear of arbitrary restraint and our aversion to the State curtailing our liberty other than as punishment for breach of the law. It is one thing when the State punishes someone for past voluntary wrongdoing; when it deprives individuals of liberty on the basis of their feared future conduct that is quite another. We resist the idea of the State taking power over individuals who have not done anything wrong and whose confinement is therefore, by hypothesis, neither deserved nor limited in its severity by the degree of their culpability. There is also the clear potential for abuse. Yet it is hard to deny that there must be some circumstances, even if of an exceptional nature, in which the State has the right to restrict liberty in order to protect the public. But if this is so, we are immediately faced with the moral problem which will be the focus of this article — the problem of ‘false alarms’ or ‘false positives’.

This problem is a function of the fact that all preventive measures are based on judgments about the likelihood or probability, not the certainty, of the individuals in question causing harm to others. Furthermore, the likelihood of feared behaviour, as I will show, is a matter of the expected frequency of that behaviour within a group to which the individual belongs. For instance, to give a highly simplified example for illustrative purposes only, suppose it is said that the likelihood of violent offenders re-offending is 0.75. This means that three out of four of the group of violent offenders will re-offend. Now let us suppose that 400 violent offenders who have served their sentences are about to be released from prison. If they are all preventively incarcerated on the basis that three in four of them will re-offend, 100 harmless individuals — false positives — will be imprisoned unnecessarily. The same problem arises when non-quantitative terminology is used. For instance, suppose it is said that violent offenders are ‘highly likely’ to re-offend. Once again, if all the members of a group of violent offenders are preventively incarcerated on the basis that they are highly likely to re-offend, some individuals will be imprisoned unnecessarily.

16 For details of the legislation, see McSherry, above n 6, 277–9. See also Ian Kerridge, Michael Lowe and John McPhee, *Ethics and Law for the Health Professions* (2nd ed, 2005) 580.

17 McSherry, above n 6, 277.

How worried about this should we be? One source of possible worry has to do with concerns about the accuracy of the risk assessments. Clearly, if the experts do not know the true risks but only claim to do so, there can be little justification for mistakenly depriving people of their liberty. But what if this concern were to be addressed? Improvements in predictive techniques might free us in the future from worries about accuracy and even in our current state of knowledge there are likely to be some circumstances in which we might expect risk assessments to be statistically sound. A deeper source of worry would nevertheless remain: namely, whether it is justifiable to deprive people of their liberty on the basis that they have a characteristic which has been proved to be statistically, though not in all cases, associated with harmful conduct. The focus of my attention will be on this more challenging question, and I will therefore assume for the sake of argument that the relevant risk assessments are statistically sound. I will, however, return in Part Five of the article to the further complications posed by the possible unreliability of the statistics.

I will discuss and give reasons for rejecting two lines of reasoning which suggest that we are entitled to downgrade any concerns we might have about the false positive problem. The first of these, discussed in Part Three, points to the fact that we tolerate the possibility of false positives in the form of wrongful convictions in the criminal context and it invites us to conclude that the possibility is no more worrying in the preventive context. In rejecting this argument, I call attention to the 'nakedly statistical' character of the generalisations on which preventive measures are based. I argue that analogously nakedly statistical evidence of guilt would be thought unacceptable because such evidence is unreliable in respect of individual defendants and is therefore in conflict with our views about the fair apportionment of the risk of error.

The second line of reasoning, discussed in Part Four, attempts to rescue preventive measures by arguing that unreliability is less of a worry in the preventive context. This line of reasoning concedes that the nakedly statistical nature of the reasoning which underlies risk assessments would be unacceptable as the basis for the imposition of criminal liability, but argues that in the case of non-punitive measures it is reasonable to balance the interests of 'risky' individuals against those of society, asking them to bear the risk of error in the interest of public safety. I argue that this view exaggerates the moral differences between punitive and protective measures. In its place, I propose a rights-based view, on which we recognise that people who present a risk of harm have the right that the risk of error be skewed in their favour. This right is of the same character and rests on the same values as the right of criminal defendants to be protected against mistaken conviction. I argue that the risk of error should be borne mainly by the community because there are certain fundamental values at stake which take priority over the interest in public protection and which therefore cannot simply be balanced against it. In Part Five of the article, I explain how my views about the fair apportionment of the risk of error in the preventive context can be translated into appropriate legal standards. Finally, in Part Six, I argue that my analysis helps us to understand the meaning of 'arbitrariness' in the context of the prohibitions on arbitrary detention and the arbitrary invasion of privacy which are contained in recent human rights legislation.

3. *Naked statistics*

Alan M Dershowitz points out that probabilistic judgments are involved not only when legal decision-makers are called upon to make prospective decisions but also when they are required to make retrospective decisions. After all, when it comes to determining whether a disputed past event occurred, absolute certainty is not the standard. It is only necessary to decide ‘beyond a reasonable doubt’ in the criminal law context and by ‘a preponderance of the evidence’ in the civil law context that a past event occurred.¹⁸ Since the use of probabilistic standards in criminal and civil law opens the door to false positives — namely, wrongful convictions and erroneous impositions of liability — Dershowitz implies that we need not be overly worried about the false positives which result from using probabilistic standards in the preventive context. Christopher Slobogin makes a similar point. He argues that the culpability determinations made in a criminal trial are ‘subject to serious inaccuracy’,¹⁹ and that ‘[i]f we are willing to countenance a criminal system based on this degree of uncertainty, we may be hard-pressed to criticize a preventive detention regime on unreliability grounds.’²⁰

I will concentrate in what follows on the criminal law context. It is obviously correct to say that legal fact-finders undertake a probabilistic inquiry. Notwithstanding the reluctance of courts to quantify the probabilities, this is an inevitable consequence of the fact that legal fact-finding takes place under conditions of uncertainty and that we allow individuals to be convicted by evidence that carries some risk of inaccuracy.²¹ I will argue, however, that there is a difference between the kinds of probabilistic judgments made in the course of a criminal trial and estimates of the probability of future harmful conduct. This argument has nothing to do with the criminal standard of proof because even predictions of future harmful conduct which can be proved beyond reasonable doubt are based on a form of reasoning which would be unacceptable in the criminal context, as we will see.

I will start with the probabilistic judgments which underlie risk assessments. There are two approaches to risk assessment: clinical and actuarial. Most of the discussion about these approaches has taken place in the context of attempts to assess the likelihood of violent recidivism and, for ease of exposition, I will concentrate on this context in what follows. However, the points I will make are points of principle and can readily be extrapolated to other contexts in which individuals are judged to be a risk to society. The two methods have been explained as follows:

One approach, called clinical prediction, relies on the subjective judgment of experienced decision makers — typically, in the case of violence, psychologists

18 Alan M Dershowitz, *Preemption: A knife that cuts both ways* (2006) 16–17.

19 Christopher Slobogin, ‘A Jurisprudence of Dangerousness’ (2003) 98 *Northwestern University Law Review* 1, 7.

20 *Ibid* 8.

21 Alex Stein, *Foundations of Evidence Law* (2005) 81: ‘When a particular occurrence is neither absolutely certain nor altogether impossible, then it is probable’.

and psychiatrists, but also parole board members or judges. The risk factors assessed in clinical prediction might vary from case to case, depending on which seem more relevant. These risk factors are then combined in an intuitive manner to generate an opinion about violence risk. The other approach, termed actuarial (or statistical) prediction, relies on explicit rules specifying which risk factors are to be measured, how those risk factors are to be scored, and how the scores are to be mathematically combined to yield an objective estimate of violence risk.²²

Actuarial violence risk assessment tools are developed empirically by following released violent offenders with a view to finding out which of them re-offend within a specified time period. The characteristics of those who re-offended are then compared with the characteristics of those who did not re-offend, with a view to identifying which characteristics are predictors of subsequent violence. These are known as the ‘predictor variables’. Once the predictor variables have been identified, they are then assigned weights, depending on how strongly they are correlated with violent recidivism. The variables are then combined to form a scale, allowing us to say that an individual with a particular score on the scale has a certain probability of re-offending — say, 75 per cent — because he or she shares characteristics with a group of offenders, 75 per cent of whom were observed over the follow-up period to re-offend. The scores on the scale, in other words, translate into quantitative estimates of risk. As Eric S Janus and Robert A Prentky explain, ‘actuarial assessment tells us the empirically measured rate of recidivism among a group of ... offenders who share a set of characteristics with the subject of the evaluation’.²³

The Violence Risk Appraisal Guide (‘VRAG’) is a well known instrument of this form. It identifies twelve factors as particularly significant in predicting violent recidivism and places individuals into one of nine categories based upon their actuarial risk of future violence. The factors are: (1) score on the Psychopathy Checklist; (2) separation from parents before age 16; (3) victim injury in index offence; (4) schizophrenia; (5) never married; (6) elementary school maladjustment; (7) female victim in index offence; (8) failure on prior conditional release; (9) property offence history; (10) age at index offence; (11) alcohol abuse history; (12) diagnosis of personality disorder.²⁴ These variables are numerically combined to form the VRAG.

If a score on an instrument such as the VRAG is used as the basis for restricting an individual’s liberty, his or her liberty is restricted on the basis of the generalisation that members of the group to which he or she belongs are more likely to commit violent acts than members of other groups. It might seem that

22 John Monahan, ‘A Jurisprudence of Risk Assessment: Forecasting harm among prisoners, predators, and patients’ (2006) 92 *Virginia Law Review* 391, 405–6.

23 Eric S Janus and Robert A Prentky, ‘Forensic Use of Actuarial Risk Assessment with Sex Offenders: Accuracy, admissibility and accountability’ (2003) 40 *American Criminal Law Review* 1443, 1476.

24 Monahan, above n 22, 410, n75. Monahan notes that subjects who injured a victim in the index offence, who were diagnosed as schizophrenic, who chose a female victim for the index offence, or who were older, were significantly less likely to be violent recidivists.

clinical judgment is different, in apparently providing a more ‘individualised’ assessment of risk, but, as Janus and Prentky show, clinical judgment also relies on group-based generalisations.

It is true that clinicians are not committed to a predetermined set of factors with fixed weights but are free to focus on the features of the individual which they think most significant. They are, however, still relying on group-based generalisations drawn from experience with other offenders with similar case histories or similar symptoms, even if these generalisations are not expressly articulated and their reliance on them may therefore not be obvious. When clinicians predict that an individual offender is, let us say, highly likely to commit a future act of violence, they must be relying, as Janus and Prentky say, on ‘perceived commonalities with similarly-situated [sic] others — i.e., comparisons to group characteristics and outcomes — ascertained by clinicians in their training and experience.’²⁵ If they were not implicitly comparing the individual with other offenders with similar characteristics of whom they have had experience — a so-called ‘reference class’, even if a personally constructed one — their judgment would be a guess, not a prediction.²⁶ Both clinicians and statisticians therefore use membership in a group (having certain traits or characteristics or fitting a certain ‘profile’) as a probabilistic (but, of course, not certain) sign of being a danger to other people. (Whose predictions are more accurate is a matter I will address briefly in Part Five of this article.)

Frederick Schauer makes a similar point in his book, *Profiles, Probabilities and Stereotypes*.²⁷ One of his examples is the use of profiles by officials such as customs officers. Consider a profile for drug couriers. Such a profile might take the form of a predetermined set of factors which are mechanically applied. For instance, customs officials might routinely subject to intensive scrutiny anyone who is coming from a country known to be a supplier of drugs, has paid cash for their ticket, is not a member of a frequent flyer programme, travels more than might be expected from their occupation, is travelling alone and is wearing loose-fitting clothes. (All of these factors are apparently probabilistically correlated with being a drug courier.) On the other hand, the officials might not use a predetermined formula but might make a case-specific decision based on their subjective best judgment as to which travellers should be subjected to close scrutiny. The unwritten, flexible and seemingly more individualised approach might appear to be different from the mechanical approach, but, as Schauer points out, there is, at most, a difference of degree here, not of kind. The former is based just as much on probabilistic generalisations drawn from previous experience and

25 Janus and Prentky, above n 23, 1477. See also Barbara D Underwood, ‘Law and the Crystal Ball: Predicting behavior with statistical inference and individualized judgment’ (1979) 88 *Yale Law Journal* 1408, 1427: ‘Although the clinician need not identify in advance the characteristics he will regard as salient, he must nevertheless evaluate the applicant on the basis of a finite number of salient characteristics, and thus, like the statistical decision-maker, he treats the applicant as a member of a class defined by those characteristics.’

26 Janus and Prentky, above n 23, 1477–9.

27 Frederick Schauer, *Profiles, Probabilities, and Stereotypes* (2003).

is therefore just another kind of profiling. Whether the profile is constructed in advance or intuitively on a case-by-case basis, it involves singling out certain individuals for examination on the basis that they share certain characteristics with people who have been found in the past to be drug couriers.²⁸

By contrast, probabilities based on nothing more than group membership play nothing like this role in finding someone guilty of an offence. Consider the following case, constructed by Charles R Nesson, a case which I will call ‘Prisoner One’.

In an enclosed yard are twenty-five identically dressed prisoners and a prison guard. The sole witness is too far away to distinguish individual features. He sees the guard, recognizable by his uniform, trip and fall, apparently knocking himself out. The prisoners huddle and argue. One breaks away from the others and goes to a shed in the corner of the yard to hide. The other twenty-four set upon the fallen guard and kill him. After the killing, the hidden prisoner emerges from the shed and mixes with the other prisoners. When the authorities later enter the yard, they find the dead guard and the twenty-five prisoners. Given these facts, twenty-four of the twenty-five are guilty of murder.²⁹

If one of the prisoners (call him ‘John’) is now charged with murder, there is a very high probability — a 96 per cent chance — that he is guilty. The evidence supporting his guilt is, however, ‘nakedly statistical’.³⁰ Alex Stein explains the concept of ‘naked statistics’ as follows. It means:

any information about a category of people or events not evidencing anything relevant in relation to any person or event individually. A piece of evidence is nakedly statistical when it applies to an individual case by affiliating that case to a general category of cases.³¹

Would John’s membership in the reference class of inmates in the yard, which makes him very likely to have committed murder, be sufficient to convict him? And would it be just to convict him on this basis? Despite the fact that the criminal standard of proof is not proof beyond doubt but proof beyond reasonable doubt, the answer to both the legal and the moral question would appear to be ‘no’. Even if the probability of John’s guilt were higher — suppose there were 100 or even 1000 prisoners in the yard, of whom one is innocent — our intuitive response would not change. Describing a similar case, albeit a civil one, James Franklin says there would be an element of randomness in the attribution of liability.³² But why is this so?

28 Ibid 167–74.

29 Charles R Nesson, ‘Reasonable Doubt and Permissive Inferences: The value of complexity’ (1979) 92 *Harvard Law Review* 1187, 1192–3.

30 This term was first used by David Kaye, ‘Naked Statistical Evidence’ (1980) 89 *Yale Law Journal* 601.

31 Stein, above n 21, 43.

32 James Franklin, ‘Case Comment – *United States v. Copeland*, 369 F. Supp. 2d 275 (EDNY 2005): Quantification of the “proof beyond reasonable doubt” standard’ (2006) 5 *Law, Probability and Risk* 159, 162.

One answer is that we resist the imposition of liability in Prisoner One because we feel that there must be some additional, case-specific evidence about some of the perpetrators — such as trace evidence or a history of hostility towards the guard — and the failure of the prosecution to adduce such evidence suggests suppression of evidence, decreasing the probability that the prosecution's case is correct.³³ But suppose — as is not impossible — that the prosecution is able satisfactorily to explain the absence of case-specific evidence. It seems that we still would not wish to convict John. We therefore need a better account of why this is so.

A different objection to the use of naked statistical data in court, discussed by Laurence Tribe, is based on the confusing nature of such data, the fact that people often disagree about its correct analysis, and its consequent potential to lead fact-finders into error.³⁴ But this also does not get to the heart of the matter because there can be reasons to object to the use of naked statistical evidence even when it is easy to apply to the case at hand, as indeed it is in the case of Prisoner One. Even in such cases, where the evidence is neither complex nor potentially misleading, we think it unjustifiable to convict an individual of a crime merely on the statistical ground that he or she belongs to a group almost all of the members of which are guilty of the crime.

The best way of accounting for our intuitive response to cases such as Prisoner One is to draw attention to the fact that numerically high probabilities are not necessarily 'weighty'. The concept of weight was first suggested by John Maynard Keynes,³⁵ and L Jonathan Cohen has subsequently elaborated upon it.³⁶ In order to understand it, we need to start with the fact that the selection of reference class can make a large difference to probability assessments.³⁷ Consider John. He is a member of the reference class of inmates in the yard, 24 of 25 of whom are guilty of murder. This means that the probability of his being guilty of murder, on the evidence that he was in the yard, is 96 per cent. Cohen describes this as 'a generalised judgment of conditional probability, in the sense that it does not assert anything about a particular person. From it we can derive its instantiation for a particular person only on the assumption that our reference to the person does not add to or subtract from the evidence on which the probability is conditional.'³⁸

This assumption cannot, however, be made since, for all the court knows, John is different from the other prisoners, or most of the other prisoners, in a respect which would lead it to downgrade its probability estimate. Suppose, for instance,

33 See Mark Kelman, 'The Necessary Myth of Objective Causation Judgments in Liberal Political Theory' (1987) 63 *Chicago-Kent Law Review* 579, 592.

34 Laurence H Tribe, 'Trial by Mathematics: Precision and ritual in the legal process' (1971) 84 *Harvard Law Review* 1329, 1334–7.

35 John Maynard Keynes, *A Treatise on Probability* (1921) 71–7.

36 L Jonathan Cohen, *The Probable and the Provable* (1977) 36–9.

37 Mark Colyvan, Helen M Regan and Scott Ferson, 'Is it a Crime to Belong to a Reference Class?' (2001) 9 *Journal of Political Philosophy* 168, 172. See also Ronald J Allen and Michael S Pardo, 'The Problematic Value of Mathematical Models of Evidence' (2007) 36 *Journal of Legal Studies* 107, 113.

38 L Jonathan Cohen, 'Twelve Questions about Keynes's Concept of Weight' (1986) 37 *British Journal for the Philosophy of Science* 263, 265–6.

to use a categorisation suggested by Cohen in the course of discussing an analogous case, that the prisoners can be divided into ex-boy-scouts and non-scouts and that the former are much less inclined to violence than the latter. John is an ex-boy-scout. If the court were to use the reference class of 'ex-boy-scouts in the yard' rather than the reference class of 'prisoners in the yard' to assess the probability of John's guilt, the probability that he is guilty would be lower than 96 per cent.³⁹ Or suppose that the prisoners can be divided into those over age 60 and those under age 60 and that prisoners over age 60 are much less likely to commit crimes of violence. If John is over the age of 60, the probability that he is guilty is once again less than 96 per cent. The point is that there are potentially many unknown reference classes to which John belongs which might affect the probability of his being guilty. Since the probability of his guilt changes depending on which reference class is chosen, this shows that mere membership in a reference class is insufficient to justify a finding of guilt.

What more do we need? It seems that we need probabilities which are not merely numerically high but are also weighty. Weight is an independent dimension of probability, according to Cohen.⁴⁰ Unlike mathematical probability, which tells us which 'conclusion is favoured by the evidence we have',⁴¹ weight 'grades how *extensive* a coverage of the *relevant* issues is achieved by this evidence.'⁴² It follows that the fewer unknowns there are — the more variables against which the probability estimate has been tested and survived — the more weighty the estimate. Conversely, the more susceptible a probability judgment is to being overturned by a small amount of new evidence, the lower its weight.⁴³ The judgment that John is highly likely to have participated in killing the guard could be radically altered by the discovery of new evidence. One piece of plausible, reliable evidence could, as Barbara Davidson and Robert Pargetter say about a similar case, reduce the probability of guilt from extremely high to near zero.⁴⁴ Furthermore, we know that such evidence, although undiscovered, does exist for one of the prisoners.⁴⁵ The high mathematical probability that John is guilty is therefore unstable or lacks weight.

To appreciate this, contrast Prisoner One with another case, which I will call 'Prisoner Two'. In Prisoner Two there is an accumulation of different pieces of direct and circumstantial evidence which yields the same 96 per cent chance that John is guilty. Perhaps an eyewitness has testified that she saw the attack on the guard and that one of the assailants looked like John. And perhaps it can be shown that John had a motive to kill the guard. If it be thought impossible to quantify the precise mathematical chances on the basis of an accumulation of such evidence,

39 Cohen, *The Probable and the Provable*, above n 36, 78.

40 Ibid 36.

41 Ibid 39.

42 Ibid (emphasis added).

43 Cohen, 'Twelve Questions about Keynes's Concept of Weight', above n 38, 276–7.

44 Barbara Davidson and Robert Pargetter, 'Guilt beyond Reasonable Doubt' (1987) 65 *Australasian Journal of Philosophy* 182, 183.

45 Ibid 184.

we can say instead that it is highly probable that John is guilty. In Prisoner Two, we would have no qualms about John's conviction. This is in part because there is much more evidence supporting the inference that he is guilty. Although, in principle, there might be missing evidence which would lead us to change our minds, the chance that such destabilising evidence exists is very slight. Here we have a probability which is not only high but also weighty. As Franklin explains:

A probability of guilt of 0.9 reached through balancing a small amount of evidence is different from a probability of 0.9 based on a mass of evidence, because the chance discovery of a new minor piece of evidence could well reduce the first to 0.7 but is unlikely to do so for the second. One might therefore be rationally less willing to condemn a defendant to a heavy sentence on a probability of 0.9 of low weight than on a probability of 0.9 of high weight.⁴⁶

Furthermore, this body of evidence goes to the causal role played by John in the events. This is what secures its status as *relevant* evidence. In this connection Judith Jarvis Thomson makes an illuminating distinction between internal and external evidence. External evidence does not stand in an explanatory relation with the proposition for which it is evidence. Internal evidence does stand in such a relation.⁴⁷ Thomson explains this difference in the context of a civil case, but I will apply what she says to the cases of Prisoner One and Two.

In Prisoner One, the proposition that John is guilty is supported by external evidence. This is the evidence that he was a member of a group of 25 and that 24 of the group's members participated in the crime. Such evidence does not stand in an explanatory relation with the proposition for which it is evidence, namely, the proposition that John is guilty. This is because the truth of the evidence does not help us causally to connect John with the attack on the guard. Being in a group of people the majority of whom have committed a crime neither causes nor is caused by criminal activity. Put otherwise, it is pure accident that the probability of John's being guilty is 96 per cent, as is evidenced by the fact that if there had been 50 prisoners in the yard, there is no reason to think that 48 of them would have participated in killing the guard.⁴⁸

By contrast, in Prisoner Two, the proposition that John is guilty is supported by internal evidence. This is evidence which is either capable of explaining his guilt or is capable of being explained by his guilt. In particular, his motive is capable of explaining his guilt and the evidence of the eyewitness is capable of being explained by his guilt. If it is true that John is guilty, this would explain why the eyewitness thought she saw him participating in the attack. Of course, she might have misidentified him. But if her eyesight has been proved to be good, and

46 Franklin, above n 32, 162.

47 Judith Jarvis Thomson, 'Remarks on Causation and Liability' (1984) 13 *Philosophy and Public Affairs* 101, 128–31. See also Judith Jarvis Thomson, 'Liability and Individualized Evidence' (1986) 49 *Law and Contemporary Problems* 199, 203–5.

48 L Jonathan Cohen puts this by saying that a probability derived from an accidental relative frequency, unlike a probability based on a causal propensity, is not counterfactualisable because the probability could be affected by adding more or different individuals to the reference class: 'Subjective Probability and the Paradox of the Gatecrasher' (1981) *Arizona State Law Journal* 627, 633–4.

visibility on the day good, and she is a disinterested witness, then the fact that John participated in the attack will be the best explanation of her testimony. And if John had a motive for killing the guard, this is capable of explaining why he participated in the attack. His motive provides a putative cause of the events. As Thomson explains, ‘internal evidence helps us to see the event we are interested in as causally embedded in a series of events, and thus as forming part of history.’⁴⁹ A final difference between the cases of Prisoner One and Prisoner Two is the logical impossibility of attempting to rebut the evidence in Prisoner One. There is a 96 per cent chance that John participated in the attack — a chance which simply has to be accepted. This is not true in the case of Prisoner Two. In Prisoner Two, the incriminating evidence is of a kind which is susceptible to being discredited.

A number of the points made above — especially those relating to the need for evidence to be weighty and to be susceptible to challenge — inform Alex Stein’s elegant arguments in *Foundations of Evidence Law*.⁵⁰ Stein describes the Lottery Paradox. An agent has a box that contains one thousand lottery tickets. The tickets are drawn from the box one by one by the participants in the lottery. The agent knows that only one ticket in the box will win the lottery. Throughout the lottery the agent has no knowledge of the number and the outcomes of the previous drawings. It seems that the agent has good reason to believe that the first ticket will not win the lottery. It also seems the agent has the same reason to believe this about all the other tickets. But this is paradoxical because, while it seems rational to believe of each individual ticket that it will be a losing ticket, the agent also knows that one of the tickets will be a winning ticket.⁵¹ The analogy between this and Prisoner One will be clear. Since it is highly probable that John committed murder, it seems that the fact-finder has reason to believe that he is guilty of murder. But the fact-finder also has the same reason to believe this of all the other prisoners. Yet the fact-finder also knows that one of the prisoners is not guilty.

Stein’s solution to the paradox is to say that statistical evidence about a class of events cannot support findings with regard to individual members of the class.⁵² Thomson makes a similar point. She says: ‘I tear up my ticket in yesterday’s lottery when I hear a radio announcement that a different ticket won. I do not tear up my ticket in yesterday’s lottery when I hear only a million more tickets had been sold that I thought had been sold.’⁵³ Applying this point to Prisoner One, we can say that while the statistical evidence in the case enables the fact-finder to form a rational belief about the *group* of prisoners, it does not support a finding about any *individual* prisoner. What, then, would the fact-finder need to form a rational belief about the guilt of a particular prisoner such as John? Stein agrees with Cohen that the evidential base supporting the probability assessment would need to be not only numerically high but also weighty.

49 Thomson, ‘Remarks on Causation and Liability’, above n 47, 131. Richard W Wright makes a similar point in ‘Causation, Responsibility, Risk, Probability, Naked Statistics, and Proof: Pruning the bramble bush by clarifying the concepts’ (1988) 73 *Iowa Law Review* 1001, 1056–7.

50 Stein, above n 21.

51 *Ibid* 67.

52 *Ibid* 83.

53 Thomson, ‘Remarks on Causation and Liability’, above n 47, 132.

For a probability to be weighty in respect of a particular individual, it must, according to Stein, be ‘evidenced’ and a probability will be evidenced only if its evidential base satisfies what he calls the ‘principle of maximal individualization’.⁵⁴ This principle requires fact-finders to base their decisions on individualised or case-specific evidence, that is, evidence which covers the factual grounds of the accusation. It also prevents them from making a finding against a litigant when the finding is not susceptible to individualised testing by the litigant.⁵⁵ Such testing, Stein explains, ‘includes cross-examination of witnesses and all other practical means for testing evidence for credibility and for obtaining new information about the case.’⁵⁶ Consider fingerprint evidence. Although fingerprint evidence rests on a statistical generalisation — namely, the generalisation that fingerprints are virtually never identical⁵⁷ — fingerprint evidence is different from naked statistical evidence in being susceptible to individualised testing by the defendant. In particular:

Fingerprints at the scene of the crime integrate with other evidence that tells how the perpetrator committed the crime. Fingerprints on the gun with which the perpetrator killed the victim integrate with the gun and, ultimately, with the wound on the victim’s body. Fingerprints on the door of the victim’s house integrate with other evidence pointing to the defendant’s presence (or non-presence) in the house. These and other case-specific interactions make fingerprint evidence susceptible to individualized testing. For example, a defendant against whom fingerprint evidence is adduced can testify about his or her alibi.⁵⁸

Stein goes on to argue that, although there are no *epistemological* reasons to prefer ‘evidential substantiation for individual events’ over a ‘calculus of chances’,⁵⁹ the two systems allocate the risk of error in different ways and reasons of *political morality* favour the former system. In Stein’s view, if defendants were to be convicted on the basis of a probability estimate the evidential base of which does not cover the specific factual allegations brought against them and which they have not had the opportunity to test, this would expose them to too high a risk of erroneous conviction. The risk would be too high because missing evidence could force radical revisions to the probability estimate. Stein writes: ‘A proposition identified as highly probable does not have much credibility if its high probability derives from a non-weighty fact-generating argument. Such propositions are too risky to rely upon.’⁶⁰ The prosecution should therefore assume the risk of erroneous acquittal in connection with non-weighty evidence.⁶¹ It is irrelevant how high the numerical probabilities are.⁶² In Stein’s view, ‘[t]he accused ... never

54 Stein, above n 21, 72.

55 Ibid.

56 Ibid.

57 Ibid 184.

58 Ibid.

59 Ibid 206.

60 Ibid 120.

61 Ibid 177.

62 Ibid 204.

assumes the risk of erroneous conviction that accompanies evidence and inference not open to individualized testing'.⁶³ This is because criminal convictions are not legitimate unless the defendant has had the opportunity to 'disassociate his individual case from the statistically dominant category'.⁶⁴ It will be clear that Stein provides normative reasons for rejecting nakedly statistical proof of guilt. Our legal system takes the view that wrongful convictions are a greater evil than erroneous acquittals. In Stein's view, it is our wish to minimise the number of wrongful convictions which explains why we refuse to draw unreliable conclusions about the liability of particular members of a group from statistics about the group.

If we now return to the preventive context, it will be obvious that the attempt to ascribe probabilities to future harmful conduct at the hands of particular individuals on the basis of group statistics which are not weighty in respect of individual members of the group is just as unreliable as the attempt to infer liability for past conduct. The fact that we are dealing with the future rather than the past makes no difference. After all, the Lottery Paradox involved the making of predictions about future individual cases from group statistics and the same paradox arises when, for instance, 100 individuals are preventively detained on the basis that there is, say, a 0.75 probability that they will re-offend.

Although it makes sense to say that the expected incidence of re-offending in the group to which the individuals belong is 0.75, strictly speaking it does not make sense to say of any particular individual that he or she has a 75 per cent risk of re-offending or that he or she is likely to re-offend. Although I have used these phrases and will continue to do so because of their currency, the fact is that 75 of the individuals in the group will certainly re-offend (the probability of their re-offending is 1) and 25 of them will certainly not re-offend (the probability of their re-offending is 0) and it is only because we are not in possession of all the relevant information that we are unable to make these finer distinctions within the group.⁶⁵ As Bernard Robertson and G A Vignaux say, probability is not 'a property of objects and processes in the real world' but 'a measure of our own uncertainty'.⁶⁶ In a world of perfect information we would have knowledge of the additional variables which correlate with being either in the subgroup of 75 or the subgroup of 25. This is the analogue of the 'missing evidence' discussed above. However, we do not live in a world of perfect information, which means that any prediction about a *particular* individual on the basis of statistics about a group to which they happen to belong will be unfounded.

Yet, whereas we can insist on doing justice to individuals in the context of a criminal trial by refusing to attribute criminal liability on the basis of non-weighty probability judgments, we do not have a similar choice in the preventive context.

63 Ibid 177.

64 Ibid 207.

65 Janus and Prentky make this point but choose not to 'delve into this philosophical quandary': above n 23, 1477.

66 Bernard Robertson and G A Vignaux, 'Probability — The Logic of the Law' (1993) 13 *Oxford Journal of Legal Studies* 457, 460.

As we have seen, risk is a function of fitting a profile which is known to be statistically associated with causing a certain kind of harm and risk assessments are therefore necessarily based on probability data about classes of people which are insensitive to relevant but unknown differences among the individuals in that class. In seeking to prevent harm, the State will therefore be forced to rely on such data, notwithstanding their inability rationally to warrant inferences about particular individuals. But, if so, we are returned to the moral problem with which we began, the problem of false positives.

I have argued that Dershowitz's strategy for minimising the severity of the problem of false positives fails, because the unreliability introduced by reliance on statistics about relative frequencies within a reference class has no analogue in criminal trials. Perhaps, though, there is another way to discount the severity of the problem. In light of the fact that we reject nakedly statistical proof of guilt in order to minimise wrongful convictions, but are forced to accept nakedly statistical predictions of future harmful behaviour, someone might seek to justify this on the basis that entirely different moral principles apply in the protective and the criminal contexts and that the apportionment of the risk of error is the opposite from that which we think fair in criminal law. I turn to explore this possibility now.

4. *Punishment versus Protection*

The difference between criminal punishment and civil or regulatory deprivation of liberty is that the former reflects moral blameworthiness deserving condemnation whereas civil law provides protection through non-condemnatory confinement or supervision of potentially dangerous people.⁶⁷ In the view of some authors, this opens the door to balancing the individual's liberty interests against public safety in a way which we would not accept in the criminal context. In the criminal context, the rate of false positives (wrongful convictions) we are prepared to accept is very low and the rate of false negatives (acquittals of the guilty) which we accept is correspondingly high. This is not a matter of achieving the best balance of benefits over costs. We would not be willing to adjust the acceptable ratio of false positives to negatives even if the gains to the many would outweigh the costs to the few who are wrongfully convicted.

Some scholars think, however, that it is reasonable to accept whatever rate of false positives would maximise general welfare when the government's purpose is the merely regulatory one of preventing danger to the public. David Woods, for instance, argues that it is justifiable to 'redistribute the risks' presented by dangerous individuals, shifting the loss from the potential victims of harmful conduct to individuals who are likely to cause harm, because the purpose is not to punish them.⁶⁸ Alexander D Brooks, in the course of discussing the US Sexually

67 Paul H Robinson, 'Foreword: The criminal-civil distinction and dangerous blameless offenders' (1993) 83 *Journal of Criminal Law and Criminology* 693, 694, 696.

68 David Woods, 'Dangerous Offenders and the Morality of Protective Sentencing' (1988) *Criminal Law Review* 424, 431–2.

Violent Predator laws which were referred to in Part One of this article, makes a related point, using the language of balancing. He writes:

The issue is one of balancing values. An argument that emphasizes false positives to the exclusion of concern about the grave harms caused to potential victims by violent sexual offenders obscures the fact that, in deciding whether to accept predictions of future sexual violence, it is necessary to strike a balance between the risk to the offender of a mistake made when we confine him if he is nondangerous and the risk that, if we release him, he will later engage in violent sexual crimes.⁶⁹

Brooks goes on to ask a rhetorical question: ‘Is it morally wrong to make such a mistake [that of confining those who are non-dangerous]? Which mistake is more harmful in its consequence to societal values?’⁷⁰ And he answers that ‘[a] mistaken decision to confine, however painful to the offender involved, is ... simply not morally equivalent to a mistaken decision to release’.⁷¹ Nigel Walker echoes this sentiment, asking, ‘Can a mistaken release, with its tragic consequences, be counted as the arithmetical equivalent of an unnecessary prolongation of detention?’⁷² and he invites us, as Brooks does, to conclude that the former is much worse than the latter. David Boerner, also in the context of defending Sexually Violent Predator laws — in particular, Washington’s statute — uses the metaphor of ‘triage’, explaining its applicability by saying that:

One facing a situation where harm is inevitable is justified in using his abilities to minimize inevitable harm. ... [T]he pain that women and children will suffer from sexual violence and the pain that those erroneously committed will also suffer will be, in the aggregate, significantly less than would have been the case had any of the other alternatives [to Washington’s Sexually Violent Predator statute] been adopted.⁷³

These authors, in my opinion, get things exactly the wrong way around: although it is true that a mistaken decision to confine is not morally equivalent to a mistaken decision to release, it is the mistaken decision to confine which is, morally speaking, the worse mistake. Let me explain why. In the punitive context, we take the view that the risk of error should be almost entirely borne by the State regardless of the increased risk to public safety. This is because we think that the mistaken conviction of an innocent person is an *injustice*, whereas we do not think that the mistaken acquittal of a guilty person does anyone an injustice, not even those who may be harmed by that person’s subsequent conduct. We put this by saying that innocent people have a *right* to be protected against wrongful conviction by the State — a right which takes priority over whatever obligations

69 Brooks, above n 10, 752

70 Ibid 753.

71 Ibid.

72 Nigel Walker, ‘Harms, Probabilities and Precautions’ (1997) 17 *Oxford Journal of Legal Studies* 611, 612.

73 David Boerner, ‘Confronting Violence: In the act and in the word’ (1992) 15 *University of Puget Sound Law Review* 525, 576–7.

the State has to save individuals from becoming victims of violence at the hands of third parties. This is connected with strongly held intuitions that the state's negative obligations not to violate rights are more important than its positive obligations to prevent harm. It is these intuitions which underlie the criminal standard of proof and the rejection of naked statistical evidence, both of which are designed to give priority to avoiding the one kind of error — the unjust error — to the maximum extent which is feasible consistent with real-world uncertainties.

The authors quoted above suggest, however, that when we move from the punitive to the regulatory context, rights are no longer at stake and Benthamite utilitarianism suddenly holds sway. The risk of error should now fall mostly, if not entirely, on the person who is thought to pose a risk of harm because the erroneous failure to confine a dangerous person imposes greater costs on society than the erroneous confinement of individuals who are not dangerous. Such individuals can, in effect, be treated as a resource for the benefit of others. Their interests in not being mistakenly confined can be traded off against the gains to society.

Why should the fact that preventive measures are not intended as punishment make such a large moral difference? The authors I have mentioned simply assume that it does. Others seem to think that it follows from the meaning of words. Slobogin, for instance, argues that different moral standards apply in the two contexts because the words 'punishment' and 'prevention' apply to different kinds of measures. He argues that procedural protections of the kind available in a criminal trial are not required in preventive detention proceedings just because the proceedings are not criminal:

Criminal punishment is based solely upon a conviction for an offense and can occur only if there is such a conviction. Preventive detention is based solely upon a prediction concerning future offenses and can occur only if there is such a prediction. Therefore, preventive detention is not criminal punishment. Indeed, the concept of 'punishment' for some future act is incoherent. Accordingly, to the extent procedural protections depend upon characterization of a proceeding as criminal, they are not required in preventive detention proceedings.⁷⁴

This argument is unconvincing. It attempts to draw a substantive moral conclusion from a fact about labels. Even if it is true that it is logically impossible to apply the label 'punishment' to protective measures,⁷⁵ it does not follow that it is morally justifiable to apportion the risk of error differently in the two contexts. There might be underlying moral similarities between confinement for punitive and protective reasons, giving rise to a right to be protected against mistaken confinement regardless of the reason for it. In what follows I will suggest that this is indeed the case. I will argue that, although in the case of some regulatory measures it is reasonable to balance the cost of false positives against the cost of false negatives,

⁷⁴ Slobogin, above n 19, 12–13.

⁷⁵ For a contrary view, see Patrick Keyzer 'Preserving Due Process or Warehousing the Undesirables: To what end the separation of judicial power of the Commonwealth?' (2008) 30 *Sydney Law Review* 101, 108–9; Denise Meyerson, 'Using Judges to Manage Risk: The case of *Thomas v Mowbray*' (2008) 36 *Federal Law Review* 209, 224.

when it comes to the drastic measures which are the subject of this article, justice and individuals' rights are at stake and we need to impose principled limits on the pursuit of the collective good of public safety.

I will begin with cases in which a balancing approach is acceptable. As Schauer points out, the distribution of burdens on the basis of statistically sound but non-universal generalisations is frequently routine and defensible. Think, for instance, of laws which prevent individuals under a certain age from driving a car or buying alcohol,⁷⁶ or which subject airline pilots to a mandatory retirement age,⁷⁷ or which place restrictions on owning a particular breed of dog which is thought to be dangerous.⁷⁸ The generalisations in question — those under a certain age are not responsible enough to drink or to drive, airline pilots over a certain age have slower reflexes and diminished hearing and vision, pit bulls are aggressive — have a sound statistical basis and although they do not apply to all members of the class, we are willing to tolerate their considerable over-inclusiveness. This is for reasons of efficiency and practicality and also because there is no hint of animosity or prejudice against those in the class — the false positives — who do not have the relevant characteristics but are nevertheless burdened by the policy in question. Although Schauer does not say this, it seems that in the cases he describes we are prepared to balance the harm done to the individuals who are unnecessarily burdened by the law against the social benefits of regulation. All we ask is that the measure in question should be a rational way of achieving a legitimate governmental interest.⁷⁹

But it is very different with the sorts of measures which are the subject of this article. Consider, for instance, the case of someone who is deprived of liberty because they have a certain score on the VRAG scale. Although the attribute of dangerousness is not possessed by all members of the class (namely, by all individuals with that score), it is nevertheless probabilistically indicated by membership in the class. (It will be remembered that I am assuming that the risk assessments which are the subject of this article are statistically sound.) The interference with liberty therefore has a rational basis and the governmental interest — in protecting public safety — is clearly legitimate. But is that sufficient to justify restricting the liberty of everyone in the class?

In my view, it is not. First, when falsely attributing a characteristic to an individual causes severe hardship, such as drastic restraints on liberty or the invasion of privacy, we are much more reluctant to use group membership as a proxy for having the characteristic, even if membership in the class is a

76 Schauer, above n 27, 120.

77 Ibid 108.

78 Ibid 56.

79 This is the 'rational basis' test in United States equal protection law. It was first articulated in the case of *Gulf, Colorado & Santa Fe Railway Co v Ellis*, 165 US 150, 151 (1897): 'The mere fact of classification is not sufficient to relieve a statute from the reach of the equality clause of the Fourteenth Amendment, and in all cases it must appear not merely that a classification has been made, but also that it is based upon some reasonable ground — something which bears a just and proper relation to the attempted classification, and is not a mere arbitrary selection.'

probabilistic indicator of having that characteristic. If young people who are responsible enough to drink or to drive are not allowed to do so, or airline pilots whose reflexes, sight and hearing are undiminished are forced to retire at a particular age, we are less likely to be concerned than if more fundamental interests such as those in liberty and privacy are mistakenly invaded. In such circumstances, we are much more worried about the fact that the applicable generalisation, although generally reliable, is not universally applicable. The fact the purpose of the measure is not punitive does not diminish our concerns.

A second reason to be concerned about non-universal generalisations, as Schauer himself concedes, is that they may strike at the value of equality. This will be so when they use categories such as race or gender as predictors of certain characteristics. We are very suspicious of such generalisations, even if race and gender are statistically associated with the characteristics in question and even if the classifications are used to serve a legitimate goal. This is because, as Schauer says, certain forms of generalisation are 'morally repugnant because of the way in which they may stigmatize or isolate members of certain traditionally oppressed or marginalized groups.'⁸⁰ Members of groups defined along lines such as race and gender have often been the victims of mistaken generalisations based on nothing more than bias and prejudice, and even statistically sound generalisations have frequently been used to mask an underlying invidious purpose. For these reasons, even when the characteristics in question are statistically relevant and used for a legitimate purpose we are much more reluctant to burden everyone in the group with a generalisation which applies to only some of its members.

Consider, for instance, the widespread objection to the use of racial profiling as a law enforcement tactic.⁸¹ Another example is provided by the National Security Regulations made during the Second World War. These authorised the indefinite detention without trial of 'enemy aliens'— a term which was used to apply to individuals who had German, Italian or Japanese origins, even if they were British subjects by birth or naturalisation. We now look back with shame on the way in which the authorities used race and ancestry as markers of potential disloyalty.⁸² And finally, consider the invidiousness of using race or ethnicity to achieve a protective purpose in the sentencing context. John Monahan makes reference to a penalty hearing of a United States capital murder case in which the jury had to decide whether the defendant, if not executed, 'would commit criminal acts of violence that would constitute a continuing threat to society.' An expert witness testified that the defendant possessed many risk factors for violence, one of which was his Hispanic ethnicity. The jury sentenced the defendant to death, a sentence which was upheld by the Texas Court of Criminal Appeals. The defendant successfully argued in the Supreme Court, however, that the use of race or ethnicity for assessing risk of future violence was, regardless of the statistical

80 Schauer, above n 27, 128.

81 For a detailed discussion of the moral problems which attach to racial profiling, see Schauer, above n 27, 175–98.

82 For discussion, see Ilma Martinuzzi O'Brien, 'Citizenship, Rights and Emergency Powers in Second World War Australia' (2007) 53 *Australian Journal of Politics and History* 207.

significance of these factors, a violation of the Equal Protection Clause.⁸³ The defendant's argument was clearly right. It is illegitimate to seek to protect society by depriving people of liberty — let alone life — on the basis that the racial or ethnic group to which they belong makes them more likely to commit violent acts. To do so is so stigmatising to members of marginalised minorities that we entirely reject the idea.

Even when the more obvious kinds of discrimination, such as racial discrimination, are not at issue, certain groups in society are at special risk for arbitrary treatment at the hands of majorities. A Note in the *Harvard Law Review* makes this point, saying that some groups are more vulnerable than others to being singled out for 'pariah status' by governments. This may be by virtue of minority status, historical discrimination, social unpopularity or political powerlessness.⁸⁴ It will be clear that the groups which tend to be targeted for coercive preventive measures are frequently vulnerable in this way. Sex offenders, the mentally ill, those suffering from diseases such as AIDS and suspected terrorists could hardly be better examples of pariah groups. This means that exaggerated predictions about their dangerousness, based on nothing more than hostility, stereotypes and fear, are likely to be too readily accepted.⁸⁵

Finally, there is a third reason to be suspicious of deprivations of liberty based on non-universal generalisations. This reason is based on the importance of respecting autonomy. Some predictions of dangerousness are not inconsistent with respect for a person's autonomy. As Barbara D Underwood points out, when 'the predicted fact is not subject to individual control, then predicting that fact is less threatening to the value of respect for autonomy. For example, prediction of violent behavior by the mentally ill ... is seldom characterized as a threat to the autonomy of the mentally ill.'⁸⁶ The same could be said about isolating someone who has a highly infectious disease, since spreading the disease is not under their control. It is very different, however, when someone is deprived of their liberty when the threat they pose is under their control. In cases such as this, preventive measures assume that people who are capable of choosing not to cause harm will cause harm, thereby denying them the opportunity to choose differently. They are treated as 'predictable objects',⁸⁷ or 'dangerous animals',⁸⁸ rather than as

83 *Saldano v State*, 70 S W 3d 873, 875 (Tex Crim App, 2002), discussed in Monahan, above n 22, 392–3.

84 Note, 'Making Outcasts out of Outlaws: The unconstitutionality of sex offender registration and criminal alien detention' (2004) 117 *Harvard Law Review* 2731, 2742, 2750.

85 For instance, a recent study which investigated the attitude of Australian legal professionals and other members of the community to involuntary treatment of the mentally ill found that both groups greatly exaggerated the likelihood of a mentally ill person being violent when compared with recent research on the topic: Judith Minster and Ann Knowles, 'Exclusion or Concern: Lawyers' and community members' perceptions of legal coercion, dangerousness and mental illness' (2006) 13 *Psychiatry, Psychology and Law* 166.

86 Underwood, above n 25, 1415–6.

87 This is Underwood's phrase: *ibid* 1414.

88 This phrase is used by Stephen J Morse, 'Blame and Danger: An essay on preventive detention' (1996) 76 *Boston University Law Review* 113, 151.

individuals with the capacity for free choice. This may be less of a worry in the case of those members of the dangerous group who would have chosen to cause harm. The threat to autonomy is, however, acute in respect of those mistakenly deprived of liberty, since they would have chosen differently.

To sum up, I have been discussing protective measures justified by non-universal but statistically sound group-based generalisations. By definition, being non-universal, some members of the group will not share the characteristics which are possessed by the majority of the group's members and the measures in question will therefore burden them mistakenly. I have argued that although such mistakes can sometimes be justified on the basis of a cost-benefit calculation, this will not always be the case. Notwithstanding the fact that the measure is regulatory not punitive, if it affects fundamental interests such as liberty or privacy, or is based on invidious classifications which tend to be an arbitrary or irrelevant basis for different treatment, or strikes at the value of free choice, balancing the costs to the mistakenly affected individuals against the public's interest in security will be illegitimate.

We should not, however, go to the other extreme and say that only people who actually *will* cause harm can be justifiably deprived of their liberty.⁸⁹ Such a view would make it all but impossible to take preventive action, since it is impossible to predict future events with absolute certainty. Even in the criminal context, we do not demand absolute certainty, since we tolerate some risk of wrongful conviction. What my discussion is intended to show is not that preventive measures are entirely unjustifiable but rather that the very strong competing individual interests that I have identified are too important to be simply subordinated to the utilitarian calculus. By contrast with 'ordinary' interests, such as owning a particular breed of dog, the interests in privacy, freedom from restraint, non-discrimination and autonomy need special protection from State interference and their invasion consequently requires much more than a legitimate State interest and a rational basis. This is to say that these more fundamental interests should take second place to public safety only in exceptional circumstances. In spelling out the nature of these exceptional circumstances, it will be useful to draw on the equal protection jurisprudence of the United States Supreme Court, which subjects different kinds of classificatory laws to different levels of scrutiny.

Consider the case of *Craig v Boren*,⁹⁰ in which the Supreme Court struck down a law which prohibited the sale of alcohol to males under the age of 21 and to females under the age of 18. The law was based on statistical evidence that males between the ages of 18 and 20 were more likely to drink and drive than women of that age. Accepting for the sake of argument that the evidence was correct, the Court nevertheless took the view that gender is a 'quasi-suspect' criterion and that gender classifications therefore require an 'important', not merely a legitimate, governmental objective. The Court also held that there must be a tighter fit than

89 Von Hirsch makes this suggestion: Andrew von Hirsch, 'Prediction of Criminal Conduct and Preventive Confinement of Convicted Persons' (1971-2) 21 *Buffalo Law Review* 717, 743, n74.

90 *Craig v Boren*, 429 US 190 (1976) ('*Craig*').

usual between the state's purpose and the criterion used: the classification must be 'substantially', not merely rationally, related to the achievement of the government's objective.⁹¹ Since there are many men under the age of 21 who would not drink and drive (and many young women who would), the Court invalidated the law as an invidious discrimination against males of 18 to 20 years of age. This seems correct. There is a moral difference between a law which prevents anyone under the age of 21 from buying alcohol and the law which was considered in *Craig*. Although the former law may be just as over-inclusive as the law considered in *Craig*, we are much more willing to tolerate the errors which attend it because generalisations based on age are less likely to be invidious than generalisations based on gender.

I have explained the Supreme Court's approach when dealing with quasi-suspect classifications, such as gender. The Court is even more demanding when dealing with classifications based on race, which it regards as 'suspect'. When someone is disadvantaged by a suspect criterion, such as race, it demands a 'compelling interest' and insists that the means be even more narrowly tailored to the end.⁹² In effect, as Owen Fiss explains, 'any degree of avoidable overinclusiveness or underinclusiveness would be deemed "too much".'⁹³ This also seems correct. The Court's approach is merely a reflection of our intuitive moral responses to racial classifications: we are extremely suspicious of such classifications and there are very few circumstances in which we are prepared to countenance their use.

I suggest that such an 'intensified scrutiny' approach is well suited to the coercive measures which are the subject of this article. It enables us to give content to the idea that our fundamental interests in freedom from restraint, non-discrimination and autonomy cannot be simply traded off against the collective good of public safety. It does so by insisting that the government's objective must be sufficiently important and that its assessment of risk be based on more than rough or approximate generalisations. If we hold government to more stringent standards of this kind, the effect will be to give priority to avoiding false positives over avoiding false negatives, thus imposing the brunt of the risk of error on society in the same way as we do in the criminal context. Even if the risk of error should not be skewed in favour of the individual to quite the same extent — the exact standards will be canvassed in Part Five below — there will be at most a difference of degree, not kind, when it comes to fairly apportioning the risk of error in the law enforcement and preventive contexts.

Furthermore, in both contexts, the explanation of why we should skew the risk of error in favour of the individual is the same: an individual's interest in not being mistakenly deprived of liberty by the State is, like their interest in not being wrongfully convicted, more fundamental than an individual's interest in not

91 *Craig*, 429 US 190, 197 (1976).

92 *Adarand Constructors, Inc v Pena, Secretary of Transport*, 515 US 200, 235 (1995).

93 Owen Fiss, 'Groups and the Equal Protection Clause' (1976) 5 *Philosophy and Public Affairs* 107, 113-4.

suffering harm at the hands of third parties whom the State has mistakenly left at large or mistakenly acquitted. It is the difference in character of the competing interests — their incommensurability or insusceptibility to being measured on the same scale — which precludes a consequentialist balancing exercise.⁹⁴ This means that the metaphor of ‘triage’ is entirely inappropriate.

Triage applies when we have to choose whom of a group of people to help or even — supposing that whatever we do someone will be harmed — whom of a group of people not to harm. Suppose, for instance, to use an example first introduced by Philippa Foot, that the driver of a runaway tram can steer only from one narrow track onto another and that five men are working on the one track and one man is working on the other. Most people will say that the driver ought to steer for the less occupied track. Similarly, as Foot also points out, if a doctor has a limited quantity of a life-saving drug and has to choose between giving it to one patient, who will die if they do not receive all of the drug, and five patients, each of whom can be saved by one-fifth of the dose, it is legitimate to choose to save the five rather than the one. In these examples, we have to decide which of our positive duties (our duties to provide aid) we should fulfil and which of our negative duties (our duties to avoid injury) we should fulfil, in circumstances in which it is impossible to fulfil all of our positive duties or all of our negative duties. It is because the duties between which we are obliged to choose are of the same kind or character that balancing or weighing is the appropriate decision-procedure.⁹⁵ By contrast, we cannot simply balance a mistaken deprivation of liberty against the gains to the public because here the conflict is between interests of a different character, the negative interest in not being mistakenly deprived of liberty being much more fundamental than, and taking priority over, the positive interest in being saved from harm.

Stein also thinks that cases in which people stand to be deprived of their civil liberties are like criminal trials, in that an asymmetrical risk of error should be faced by the parties.⁹⁶ His reasoning is, however, different. He thinks that the asymmetry in both cases can be explained in utilitarian terms. Thus the very high criminal standard of proof is explained, in his view, by the fact that convicting an innocent person is much more socially harmful than acquitting a guilty person. Stein thinks, indeed, that the standard of proof can be expressed with mathematical exactitude once we know the ‘disutility differential’ between wrongful convictions and wrongful acquittals. Consider Blackstone’s maxim that ‘it is better that ten guilty persons escape, than that one innocent suffer.’⁹⁷ If true, this should lead us, Stein argues, to set the criminal standard of proof at a particular level. In particular,

94 For discussion of the idea that values are not necessarily commensurable, see Frederick Schauer, ‘Commensurability and its Constitutional Consequences’ (1993–4) 45 *Hastings Law Journal* 785; Denise Meyerson, ‘Why Courts should not Balance Rights against the Public Interest’ (2007) 31 *Melbourne University Law Review* 873, 887.

95 Philippa Foot, ‘The Problem of Abortion and the Doctrine of the Double Effect’ in Philippa Foot (ed), *Virtues and Vices, and Other Essays in Moral Philosophy* (2002) 19, 23–4.

96 Stein, above n 21, 153.

97 William Blackstone, *Commentaries on the Laws of England* (1st ed, 1769) vol IV, 352.

adjudicators should convict when they believe that the probability of a defendant's guilt is greater than 0.9. If, on the other hand, it is better from the perspective of social utility that 1000 guilty persons go free than that one innocent person be convicted, then adjudicators should convict when they believe that the probability of a defendant's guilt is greater than 0.999.⁹⁸ The criminal standard of proof, in other words, should be set by reference to whatever ratio of wrongful acquittals to wrongful convictions maximises utility. Stein writes: 'The probability threshold for convictions can thus be determined by the disutilities deriving from the socially desirable ratio of wrongful acquittals vs. wrongful convictions.'⁹⁹ Likewise, in the civil liberties context, Stein thinks that a mistaken decision to confine causes more harm on average than a mistaken decision not to confine and that it is the disutility differential between these mistakes which explains why the risk of error should be skewed in favour of the individual who poses a risk to society.

It will be evident that this is not my view. Stein treats the costs of false positives and false negatives as though they were commensurable and he therefore explains in quantitative terms why there should be an asymmetry in the risk of error: false positives cause more disutility than false negatives and that is why we should design our fact-finding and risk assessment processes to provide stringent protections against the former kind of error. Stein's reasoning is, like Brooks's, that of balancing. The only difference between them is that they disagree about which kind of mistake is more socially harmful. If Stein were to come to agree with Brooks on that matter, he would no longer think it desirable to give priority to avoiding false positives and his views about the fair way to apportion the risk of error would also change.

However, as I have shown, even if a process which tolerated more wrongful convictions and more mistaken deprivations of liberty would bring gains to society which outweigh the costs to the affected individuals — as perhaps it would — it would be wrong to apportion the risk of error differently. This is because wrongful convictions and mistaken deprivations of liberty are not just another bad consequence to be weighed against the bad consequences of releasing a guilty or a dangerous person. They are errors of a different *kind* — unjust errors — not merely errors with a negative utilitarian value: rights and justice are at stake, not merely aggregate utility.

5. *How to Minimise False Positives*

I turn now to the precise nature of the protections which are needed if the community should bear the brunt of the risk of error when the State seeks to deprive people of their liberty for preventive reasons. At the outset, it is necessary to point out that there are two different probability standards which need to be set at an appropriate level. First, we need to set a standard for the probability of the feared behaviour itself. Secondly, we need to set a standard for the probability that the method for determining the probability of the feared behaviour is reliable.

98 Stein, above n 21, 148–9.

99 Ibid 172.

Grant H Morris makes a similar point. He talks of a substantive standard and a procedural burden. The substantive standard relates to how probable the harm must be before the coercive measure is justified. In the hypothetical example with which I began this article, there was a 75 per cent probability of harm. Should we be satisfied with this or should we demand a higher probability? The procedural burden relates to provability. How much confidence should we have in the prediction before the measure is justified? Should it, for instance, be proved beyond reasonable doubt or by a less demanding standard?¹⁰⁰

The procedural standard raises the matter of the reliability of the risk assessments. We are certain that the probability of a coin coming up heads is 0.5 and so far I have assumed that we are equally certain about the probabilities which are the focus of this article. This was, of course, an idealisation. It is very rare to know probabilities with this kind of precision. How justified are we, then, in relying on the probabilistic judgments which are used to deprive people of liberty for purposes of social control? How much confidence, in other words, do we really have in the experts' claims about the risks? If the experts are prone to exaggerate the risks, ascribing high probabilities when the probabilities are in reality much lower, this would obviously aggravate the false positive problem and, given sufficient inaccuracy, one would have to question whether confinement of the individuals who are the subject of the predictions even has a rational connection with the goal of public safety.

The issue of reliability has been investigated in relation to the clinical and actuarial predictions of future violence which were discussed in Part Three of this article. The consensus appears to be that the actuarial approach is superior to clinical judgment.¹⁰¹ It seems that psychiatrists' and psychologists' predictions of future dangerousness are extremely unreliable and that they are very prone to 'over-prediction' or excessive caution. The figure usually cited is that for every three persons predicted to be dangerous by mental health professionals, only one actually re-offends.¹⁰² Of course, even if the actuarial predictions are less hazardous than clinical predictions, this does not mean that they are reliable enough: questions have also been asked about the validity of these instruments and about how good the science is which lies behind them, and it is generally accepted that their shortcomings detract from their reliability.¹⁰³ In general, there are many studies which focus on the methodological and conceptual difficulties in attempting to assess the risk of violence.¹⁰⁴ The accuracy of these (and other) methods of risk assessment is obviously a complex empirical matter which it is impossible to resolve here, although it seems clear that there are serious problems

100 Grant H Morris, 'Defining Dangerousness: Risking a dangerous definition' (1999) 10 *Journal of Contemporary Legal Issues* 61, 77.

101 For discussion, see Monahan, above n 22, 406–8.

102 Ibid 406–7. For a more sanguine view about the accuracy of clinical predictions, particularly in so far as they are made about mentally abnormal violent sexual offenders, see Brooks, above n 10, 740–9.

103 Janus and Prentky, above n 23, 1471–4.

104 See, for instance, the discussion in Michael A Norko and Madelon V Baranoski, 'The Prediction of Violence; Detection of Dangerousness' (2008) 8 *Brief Treatment and Crisis Intervention* 73.

with some of the methods for predicting future dangerousness which have been described in this article.

It will be clear that the risk of error will be affected by both the substantive standard and the procedural burden. The lower we require the likelihood of the harm to be and the lower we require our degree of confidence in the predictions to be, the higher the risk of the erroneous deprivation of liberty. Morris's own suggestion, in the context of preventive confinement of allegedly dangerous mentally ill persons, is that there should be a 90 per cent probability of violence, suicide or self-inflicted mayhem which will occur within six months and that the State should have the burden of proving beyond a reasonable doubt that such a 90 per cent probability exists.

Morris's standards are exacting. Standards as stringent as his are likely to be easiest to meet in the public health context — for instance, when individuals who have been exposed to a very contagious and serious disease are quarantined for a short period of time on the basis of scientific evidence the validity of which is not in doubt. On the other hand, it could be argued that anything short of such stringent standards would be inappropriate, given the fact that forcing someone to give up their liberty for no reason other than the common good is such a serious step. But even if Morris's standards be thought too stringent, it is worth pointing out that there is some room to relax them while still insisting that preventive restrictions of liberty have to meet a much more demanding standard than the balancing or rational basis standard.

In so far as the procedural burden is concerned, it would not, for instance, be sufficient to prove the prediction on the balance of probabilities, because setting the procedural burden at that level would impose an equal risk of error on the person who poses a risk of harm and the community. Imposing such a symmetrical risk of error could only be justified if a false negative were as regrettable as a false positive, but I have argued that this is not the case. It follows that the State should bear the burden of persuasion at the very least by 'clear and convincing evidence'. Such a standard would let through more mistakes than the 'beyond reasonable doubt' standard but, when combined with demanding substantive requirements, the risk of the mistaken deprivation of liberty would be diminished. In particular, the person should be not merely likely but very likely to cause harm, the harm should be grave, and the risk of it should be in the near future.¹⁰⁵ Such requirements would, no doubt, be met only in very exceptional circumstances, but that is as it should be, given the fact that rights are at stake, not merely aggregate utility.

105 Although the Sexually Violent Predator Laws in the United States generally require proof beyond reasonable doubt, the procedural protection is offset by very weak substantive criteria for commitment, greatly increasing the chance of false positives. This confirms the need to set both the substantive and the procedural standard at an appropriately high level. For a comparative analysis of the US standards, see Adult Diagnostic and Treatment Center, Inmate Resident Committee, Legal Subcommittee, 'Inside Civil Commitment: Competing rights, competing interests' (2003) 13(1) *Issues in Child Abuse Accusations*, app B <http://www.ipt-forensics.com/journal/volume13/j13_1_3.htm> at 20 March 2009.

There would, of course, have to be additional protections against erroneous deprivation of liberty, such as a full hearing, with access to a legal representative and the right to present evidence and cross-examine witnesses, and periodic review of the continuing lawfulness of the order. I will not fill in the details here because I am more concerned to make the point of principle, namely, that when the community seeks to curtail a person's liberty for no reason other than its own benefit, it should bear most of the risk of error.

6. *Human Rights Legislation and Preventive Measures*

My discussion so far has been entirely from the perspective of political morality. I have not discussed the possibility of legal challenges to laws depriving individuals of liberty for preventive reasons. Such a discussion would have to take in the compatibility of the relevant laws with the separation of judicial power as contained in the *Commonwealth Constitution*. It would also have to cover possible challenges under domestic human rights legislation – the *Human Rights Act 2004* (ACT) and the *Charter of Human Rights and Responsibilities Act 2006* (Vic). I will set aside the former issue, as its dimensions are well known, and end with some brief comments about the latter.

Although there are a number of provisions in the ACT and Victorian human rights legislation which could potentially have a bearing on the laws which are the subject of this article, I will consider only two of them. These are the provisions conferring the right not to be subjected to arbitrary detention¹⁰⁶ and the right not to have one's privacy, family, home or correspondence interfered with unlawfully or arbitrarily.¹⁰⁷ But when are detention and the invasion of privacy 'arbitrary'? My analysis can help to answer this question.

This article has dealt with the moral problem of false positives: individuals who are erroneously caught by over-inclusive laws. Although I have not used the language of arbitrariness, there is obviously a point at which over-inclusiveness becomes arbitrary, since measures which track the state's protective purposes too crudely or imperfectly can be described as 'arbitrary'. My analysis suggests a way of discerning when that point has been reached. The discussion in Part Four showed that the degree of imperfection or over-inclusiveness we should be willing to tolerate depends in part on what kind of interest has been curtailed. I argued that it is relatively easy for a measure which unnecessarily burdens 'ordinary' interests to be defended. Consider, for instance, a law which makes it an offence to have a blood alcohol reading over a certain level. This erroneously burdens those who can safely drive with that amount of alcohol in their blood but their interest in being allowed to do so is not of particular importance. I argued that such a law should therefore have to meet only a low or balancing standard: the state's purpose should be legitimate and the measure should be a rational way of achieving the purpose.

106 *Human Rights Act 2004* (ACT) s 18(1); *Charter of Human Rights and Responsibilities 2006* (Vic) s 21(2).

107 *Human Rights Act 2004* (ACT) s 12(a); *Charter of Human Rights and Responsibilities 2006* (Vic) s 13(a).

This low standard would be met if someone who drives with that amount of alcohol in the blood poses a more than average risk of danger on the roads. If so, the over-inclusiveness is defensible or, as we can also say, not arbitrary.

But the interests which are protected against arbitrary invasion by the ACT and Victorian bills of rights — the interests in liberty and privacy — are of a more fundamental kind and I have argued that over-inclusiveness is more worrying when fundamental interests are at stake. Such interests deserve special protection, as their inclusion in human rights legislation confirms. They should be treated as the norm and interference with them as the exception.¹⁰⁸ Interference with them therefore cannot be justified by reference to ‘loose-fitting generalities concerning the ... tendencies of aggregate groups’.¹⁰⁹ Instead, laws restricting such interests, such as preventive detention laws and laws which provide for post-sentence supervision of sex offenders in the community, need to meet less crude standards of the kind argued for in this article. In particular, the criterion on the basis of which a person is detained or their liberty invaded must be tightly, not merely rationally, connected with the state’s purpose of averting harm. Furthermore, the state’s purpose must be more than merely legitimate: the harm it aims to ward off must be both grave and imminent. If these more fine-grained standards are not met, it is legitimate to describe the detention or interference as an arbitrary exercise of power. The arguments of this article further suggest that detention and invasions of privacy will almost certainly be arbitrary if suspect characteristics such as race and ethnicity are used as predictors of future dangerousness. This is because of the arbitrariness associated with the use of racial classifications in the past, the stigmatic harm they cause, and their tendency to reinforce prejudiced attitudes and patterns of discrimination and disadvantage.

108 In the case of *United States v Salerno*, the United States Supreme Court stated that ‘[i]n our society liberty is the norm, and detention prior to trial or without trial is the carefully limited exception’: 481 US 739, 755 (1987) (Rehnquist CJ).

109 This phrase was used by Brennan J in *Craig* in relation to the ‘drinking tendencies’ of males between the ages of 18 and 20: 429 US 190, 209 (1976).

