3

Moral Responsibility Scepticism, Epistemic Considerations and Responsibility for Health

Elizabeth Shaw

3.1 Introduction

It has become increasingly common within the United Kingdom's National Health Service (NHS) for patients who have engaged in certain types of unhealthy behaviour to face restrictions on their treatment. Such policies have attracted considerable academic attention. The idea that people should be held responsible for their bad health decisions is often associated with "luck egalitarianism" (see, e.g., Albertsen and Knight 2015). Broadly speaking, this is the idea that inequalities that are purely down to luck are unjust, but that inequalities that reflect responsible agents' choices can be justifiable (Lippert-Rasmussen 2016). This chapter will discuss some difficulties faced by luck egalitarians (and proponents of similarly responsibility-sensitive approaches) in the context of responsibility for health, highlighting the implications of moral responsibility scepticism in this context. Theorists who have discussed the practical implications of moral responsibility scepticism have focused on criminal punishment. This chapter instead examines the implications of moral responsibility scepticism for whether patients should face penalties for unhealthy lifestyle choices.[1] When discussing punishment, moral responsibility sceptics often invoke an epistemic argument, maintaining that there are at least serious doubts about whether people are morally responsible (in the sense required for retributive punishment) and that, in view of this uncertainty, retributive punishment is unjust, given the serious harm it inflicts on offenders.[2] This chapter argues that this type of reasoning also implies that we

[1] I am grateful to Gabriel De Marco and Ben Davies for their helpful comments on an earlier draft of this chapter.

One of the few discussions of this topic from the perspective of moral responsibility scepticism is Levy (2018).

[2] Retributivism is the view that punishment is justified because offenders deserve to suffer hardship in return for their morally wrongful acts. There are different varieties of retributivism, cf. Moore (1997) and Duff (2001). For an overview see Walen (2020).

Elizabeth Shaw, *Moral Responsibility Scepticism, Epistemic Considerations and Responsibility for Health* In: *Responsibility and Healthcare*. Edited by: Benjamin Davies, Gabriel De Marco, Neil Levy, Julian Savulescu, Oxford University Press. © Edited by Benjamin Davies, Gabriel De Marco, Neil Levy, Julian Savulescu 2024. DOI: 10.1093/oso/9780192872234.003.0004

should not take patients' responsibility for their poor health into account when deciding whether to give these patients treatments and that the health system should not impose significant penalties on individuals for harming their own health.

After first briefly describing examples of responsibility-sensitive health care policies, Section 3.3 explains the connection between such policies and luck egalitarianism and will outline some criticisms that have been made of this view. Section 3.4 argues that (prima facie) culpability-based desert stands a better chance of helping luck egalitarians (and others with related views) to justify responsibility-sensitive health policies than some alternative approaches. However, Sections 3.5, 3.6 and 3.7 cast doubt on the idea that responsibility-sensitive health care policies based on this kind of desert are justifiable. Section 3.5 indicates reasons for doubting whether people are morally responsible for harming their own health. Section 3.6 argues that there are reasons for doubting whether harming one's own health is morally wrong. Section 3.7 argues that there are reasons for doubting that significant penalties for harming one's health would be proportionate (although very minor burdens might be less open to challenge).

3.2 Examples of Responsibility-Sensitive Health Care Policies

Such "responsibilising" policies include limiting access to treatments or deprioritising treatments for patients who are judged to be at fault for damaging their health; punishing people for poor health decisions by making them contribute financially to the costs of their treatment; and communicating stigmatic messages in health promotion campaigns about people whose illnesses are judged to be due to "bad" lifestyle choices. In England, policies requiring smokers and obese patients to take part in weight loss or smoking cessation programmes prior to elective surgery were increasingly adopted by Clinical Commissioning Groups (CCGs), organisations within the National Health Service (NHS) that made decisions about the provision of services to local populations.[3] Over a third of CCGs restricted access to some form of elective surgery for smokers or obese patients until they completed a fitness programme (Pillutla et al. 2018). The types of restriction vary (for an example of one such policy see East and North Hertfordshire CCG 2017). One justification for such policies is the forward-looking idea that patients will not sufficiently benefit from surgery until they stop smoking or lose weight. That justification has been criticised elsewhere (Royal College of Surgeons 2016). The current chapter focuses on

[3] CCGs were established under the Health and Social Care Act 2012. There were 106 CCGs in England prior to their abolition by the Health and Social Care Act 2022, which caused them to be subsumed into Integrated Care Systems.

whether responsibilising policies could be justified with reference to the backward-looking consideration that the relevant patients were responsible for engaging in behaviour that risked harming their health. However, it will suggest that other types of policy, such as taxes on cigarettes and minimum alcohol pricing, are not as open to challenge. My main aim is to discuss the moral rationale underlying responsibility-sensitivity in health care, rather than to examine the details of specific policies. Like Levy (2018), my conclusions will therefore be largely negative, that is, identifying which policy-rationales are most open to challenge rather than setting out positive policy proposals. Discussion of these underlying moral rationales is a precondition for defensible law and policy. As Levy (2018: 460) writes, "effective policy development requires the expertise of many different kinds of people with different kinds of disciplinary backgrounds. Offering purely negative advice can be an important contribution to such development."

3.3 Responsibility-Sensitive Health Care Policies and Luck Egalitarianism

The idea that people should be held responsible for their bad health decisions is part of a broader debate about whether (or how far) principles of distributive justice should be sensitive to considerations of personal responsibility (Davies 2022: 113). Proponents of responsibility-sensitivity often endorse "luck egalitarianism" (Albertsen and Knight 2015). In general, luck egalitarians claim that inequalities that are purely down to luck (i.e., circumstances over which people lack control) are unjust, but that inequalities that reflect choices for which people are responsible can be justifiable (Lippert-Rasmussen 2016). Luck egalitarians often appeal to the widespread intuition that it is fair for people to suffer or benefit from the foreseeable consequences of their own choices and unfair to ask others to bear the costs of these choices (Abad 2011), an intuition reminiscent of the saying "he's made his bed, so he should lie in it". This chapter discusses some difficulties faced by luck egalitarianism (and other responsibility-sensitive approaches) in the context of responsibility for health, focusing on the implications of free-will scepticism.

The idea that the person who is "responsible" for an action must "take responsibility" for the results rather than expecting others to do so may seem prima facie plausible, but it arguably equivocates between two senses of responsibility: "outcome responsibility" and "liability–responsibility". Drawing on the work of Hart (1968: 210–37), Vincent (2009: 46) explains that outcome responsibility is a backward-looking concept describing *past* actions as attributable to an agent. In contrast, a claim about liability responsibility makes a "prescription for the *future*" concerning what an agent must now do (or undergo) to "take responsibility" for

what has happened.[4] Applying these concepts, a more precise restatement of the luck egalitarian claim is: "if you are outcome responsible for something then you (and not others) should take liability responsibility for it" (Vincent 2009: 46). However, Vincent argues that this restatement reveals that the consequent (referring to liability responsibility) does not follow automatically from the antecedent (referring to outcome responsibility). For the conclusion about liability to follow logically from premises about outcome responsibility, one would need to add "normative bridging premises". To bridge the gap, Vincent (2009: 47) argues, we would need "reactive norms", i.e., norms that "govern our reactions to outcome responsible parties". Reactive norms are familiar from the context of criminal liability: to bridge the gap between statements such as "X is *outcome responsible* for crime Y" and "X is justifiably *liable* to receive punishment Z", utilitarian norms might be invoked that appeal to the deterrent value of imposing punishment Z on those who are outcome responsible for crime Y. An alternative reactive norm might be the retributive claim that X *deserves* punishment Z. Vincent claims that utilitarian and retributive reactive norms are also likely to feature (implicitly or explicitly) in domains beyond criminal justice when attempting to link outcome responsibility with liability responsibility. This chapter discusses problems with trying to use desert-based arguments to bolster claims that people should "take responsibility" for their own health.[5]

Examples lend intuitive support to the argument that conclusions about liability responsibility do not follow *automatically* from claims about outcome responsibility. Consider Bob. On one occasion, Bob imprudently decides not to fasten his seatbelt; his car skids on an icy road and he hits a wall. Bob sustains severe brain injuries, which he would have avoided if he had been wearing a seatbelt. Bob will die unless he receives health care. Giving him access to public resources would seem to go against the luck egalitarian maxim, because giving him these resources would require *others*—taxpayers—to "take (liability) responsibility" for the consequences of an action for which *Bob* was (outcome) responsible.[6] However, this case raises questions about the logic of luck egalitarianism, because even if we agree that Bob is outcome responsible for his injuries, due to his decision not to wear a seatbelt, this does not entail that that he should bear liability responsibility for all the consequences and be left to die. Examples like this are used to illustrate the "harshness objection" against luck egalitarianism: the objection that it is sometimes too harsh to force people to suffer all the consequences of their responsible actions (see Gabriel De Marco's chapter in this volume). Such examples also highlight the logical gap between outcome responsibility and

[4] Emphasis added.
[5] Utilitarian arguments for responsibility-sensitive health policies are not addressed.
[6] Luck egalitarians might respond in different ways to this example depending on the version of luck egalitarianism that they endorse and how they respond to the harshness objection.

liability responsibility. Outcome responsibility does not seem capable of bearing the weight that the luck egalitarian wants to put on it by itself.

Luck egalitarians might try to fill in the missing premise(s) in a way that could both (a) bridge the gap between outcome responsibility and liability responsibility, and (b) enable luck egalitarianism to avoid excessively "harsh" outcomes, such as leaving Bob to die. It might be thought that the missing premise(s) should feature the concept of "desert". In fact, Segall (2015: 355) observes that "many luck egalitarians invoke desert, whether explicitly or implicitly". As I discuss below, desert is invoked by other responsibility-sensitive distributive principles besides luck egalitarianism. I focus on a concept of desert that I call "culpability-based desert".

3.4 Luck Egalitarianism, Culpability-Based Desert and Harming One's Health

Culpability-based desert functions in the following way. To say that someone *deserves* to bear burdensome consequences[7] in virtue of having performed a risky act/omission requires meeting at least three conditions: (1) that the person was morally responsible for performing the risky act/omission, (2) that the act/omission was morally wrong and (3) that the burdensome consequences are proportionate to person's moral responsibility and the gravity of the wrongdoing.[8] My description of culpability-based desert is similar to Eyal's account of the relationship between culpable risk-taking and disadvantages, where "'Culpable' choice is understood as a free and at least somewhat morally wrong choice" (Eyal 2006: 6).

I focus on the culpability-based conception of desert for two reasons.[9] First, examples involving burdens being placed on individuals because they meet the requirements of culpability-based desert are seen by philosophers and lay people as central rather than borderline cases of "desert" (Levy 2018; Cushman 2008). Competent language-users could reasonably disagree about whether borderline cases really are examples of "desert"; or else they could agree that such examples

[7] "Burdensome consequences", or "disadvantages" to use Eyal's (2006) terminology, could refer to the natural consequences of one's act (e.g., damage to one's health) or a burden imposed by society (e.g., being made to contribute to the costs of one's treatment or being made to wait longer for treatment) in lieu of having to suffer the full natural consequences of the act.

[8] This chapter focuses on deserving burdensome consequences, rather than deserving or failing to deserve benefits. One might argue that being denied (on desert grounds) access to resources above a certain level is a case of "failing to deserve a benefit" rather than deserving a burden and that one does not need to have committed a morally wrongful act, to "fail to deserve a benefit"; perhaps failing to perform a praiseworthy act would be enough. This chapter, assumes that health care is so fundamental that the disadvantages imposed by responsibility-sensitive health care policies should be considered burdens rather than failures to benefit.

[9] For an overview of different accounts of desert see: Feldman and Skow (2020).

only involve "desert" in a loose sense.[10] In central cases, the concept of desert is used in a way that conforms with a core meaning of the term.[11] Culpability-based desert is not considered *a loose sense* of "desert" even by its critics, such as moral responsibility sceptics. Moral responsibility sceptics accept that examples of culpability-based desert involve what we standardly mean by "desert". They just argue that, given certain facts about the world, human beings cannot meet the requirements for this type of desert. A second reason for focusing on culpability-based desert is that this concept is a relatively plausible candidate (compared to some other conceptions of desert) for a principle that could both (1) bridge the gap between outcome responsibility and liability responsibility and (2) exclude harsh cases. Ultimately, however, I cast doubt on whether responsibility-sensitive health care policies based on this kind of desert are justifiable.

Some theorists who discuss desert in the context of luck egalitarianism emphasise that desert need not involve moral responsibility. Abad (2011) and Brouwer and Mulligan (2019) give the example of the most beautiful person *deserving* to win a beauty contest. This is only a borderline case of desert, involving desert in a loose sense; although it still shares some formal features with culpability-based desert. For example, there is a "fittingness" relationship between the desert-object (winning the contest) and the desert-basis (beauty); for culpability-based desert, there is a "fittingness" relationship between the desert-object (the burdensome consequences) and the desert-basis (being morally responsible for a wrongful act), which in the context of culpability-based desert could be explained in terms of proportionality.[12] The example of the beauty-contest winner only involves desert in a loose sense of the word because competent language-users disagree about whether it is a correct usage of the word "desert" to say that the most beautiful contestant deserves to win. While some might argue that the beauty-contest winner deserves to win because she satisfied the competition criteria of the competition, others might insist that, it is a misuse of the word to say that she "deserves" it, because, for example, she was just lucky to look that way. While Abad (2011) and Brouwer and Mulligan (2019) consider the case of the beauty-contest winner to be an example of desert, many writers would not, as, "by many accounts, desert seems to require agency, or the capacity to control one's own actions" (Ristroph 2006: 1301).[13] It becomes clearer that the beauty-contest

[10] There are different analyses of "borderline cases" Sorensen (2022). My description of borderline cases in terms of reasonable disagreements between language users is similar to Parikh (1994).

[11] For further discussion of the "core meaning" of terms, see Lasersohn (1999). Sometimes one word may have multiple senses, with each sense having a core meaning and borderline cases.

[12] For a discussion of such formal features of desert see Feinberg (1970: 221–50).

[13] It might be objected that such accounts are too narrow. I am not insisting that agency/control is a necessary condition for all senses of the word "desert". Rather, I maintain that generally, there is likely to be more agreement that someone "deserves" an outcome in examples where they exercise agency/control over the desert-base than in cases where they don't, and that the former are more central cases of desert. I also maintain that accounts of desert that require agency/control are of most relevance to desert-based rationales for health care policies that impose severe consequences on individuals.

example involves a loose sense of desert when the example is redescribed in a way that does not emphasise a label, such as "most beautiful", but instead emphasises the distribution of significant resources (which is the context most relevant to the current discussion). If we assume that competitions that award resources based on attractiveness are unobjectionable, it still seems to involve an odd word-choice to say that Zebedee *deserves* to receive a £1,000,000 prize because he had the genetic good fortune to have a conventionally attractive chin, whereas Xavier with his unconventional chin *deserves* to walk away penniless. I suggest that other supposedly morally neutral uses of the term "desert", such as an athlete deserving a medal, come closer than the beauty-contest example to a core meaning of desert because they have features that resemble being morally responsible (voluntarily choosing to train hard) for morally good actions (sport is assumed to be a valuable activity).

If "desert" is to provide a neat solution to the luck egalitarian's problems, the relevant sense of "desert" should help both (1) to rule out harsh cases and (2) bridge the gap between outcome responsibility and liability responsibility. Some senses of desert can do one job but not both. For example, "personhood-based desert" (Vilhauer 2013) can rule out harsh cases but cannot link outcome responsibility and liability responsibility. Personhood-based desert refers to what someone deserves in virtue of being a person, covering a diverse range of entitlements such as the right to fresh water or to receive a fair trial. We cannot forfeit through our actions things that we deserve based on our personhood. In contrast, action-based desert refers to what people deserve in virtue of their actions and omissions, not their status as persons. The concept of personhood-based desert might help explain why no one should be denied treatment for their life-threatening injuries: everyone deserves basic health care in virtue of being a person. Thus, personhood-based desert might rule out harsh cases. However, personhood-based desert cannot support the claim that people must take liability responsibility for the outcomes of their *actions*, because this claim invokes "action-based desert". Culpability-based desert is a type of action-based desert.

Other types of action-based desert (distinct from culpability-based desert) might link outcome responsibility and liability responsibility but cannot rule out harsh cases. According to one account, desert is defined in relation to social contribution (Brouwer and Mulligan 2019).[14] However, if this sense of desert is truly distinct from and does not implicitly rely on culpability-based desert, then it cannot rule out *all* cases which seem harsh (at least partly) for forcing people to suffer burdensome consequences when they were not morally responsible for running wrongful risks, or when running the risk was morally permissible or even praiseworthy, or when the burdensome consequences were disproportionate to their wrongdoing.

[14] Brouwer and Mulligan (2019: 1) argue that "luck egalitarians should consider supplementing their theory with desert considerations. Or, even better, consider desertism as a superior alternative to their theory."

It might be objected that desert does not need to serve both as a way of linking the two types of responsibility and as a way of ruling out harsh cases. Perhaps if desert just served the "linking" function, another principle could perform the function of ruling out harsh cases. Indeed, some luck egalitarians are pluralists who supplement their theory with a range of other values, such as beneficence, or concern for the worst off (e.g., Arneson 2000) which might be invoked to explain why people should not be forced to suffer severe consequences of their non-culpable choices. In reply, first, as Brouwer and Mulligan (2019: 2283) argue, "Occam's Razor tells us that we should not multiply entities without need". If we can find a concept of desert that performs both functions, that is an advantage. Second, Albertsen and Nielsen (2020) suggest that Anderson's (1999) claim that luck egalitarianism is internally inconsistent poses a problem for this pluralist solution. According to Anderson, egalitarianism assumes that each citizen should be treated with equal concern and respect, but a view that implies that someone who chose to run certain risks should be left to sink "to the depths" of severe misfortune fails to treat that person with equal concern and respect. Thus, if luck egalitarianism implies that people should suffer such harsh consequences, it cannot be rescued by invoking some other principle within a pluralist framework, because it has contradicted its initial claim to be an *egalitarian* theory. I would add that typically, in situations when one moral consideration is outweighed by another when making all-things-considered judgements, the initial moral consideration should have some intuitive force. We should feel the intuitions pulling in different directions but recognise that one ultimately has greater weight. However, I suggest there is no intuitive force to the initial claim that someone should be left to suffer severe misfortune because they non-culpably chose to run certain risks.

In contrast to these other senses of desert, culpability-based desert seems prima facie to have potential both to help justify linking outcome responsibility and liability responsibility, and to rule out harsh cases. The luck egalitarian idea that people should take liability responsibility for the consequences of actions for which they are outcome responsible seems to gain stronger intuitive support when we consider examples involving culpability-based desert, compared to cases where the prerequisites for this sense of desert are not met. Compare the car-crash victim, Bob (who merely failed to fasten his seatbelt and as a result sustained brain injuries and who does not meet the preconditions for culpability-based desert) with Celandine, who arguably meets those preconditions. Imagine that Celandine was fully morally responsible[15] for trying to run over an innocent

[15] There are different accounts of moral responsibility with different prerequisites, e.g., libertarian accounts which require indeterminism (e.g., Kane 1998) and compatibilist accounts which do not (e.g., Fischer and Ravizza 1998) and some theorists deny that anyone is ever morally responsible (e.g., Pereboom 2001; Caruso 2021a, 2021b). For the purposes of the example, I am relying on the idea that *if*

person; in her attempt to do this, Celandine crashed into a wall and sustained severe injuries. Although controversial, it seems more intuitively plausible to make Celandine bear (some of) the consequences of her actions than to make Bob do so (Schneiderman and Jecker 1996; West et al. 2003; Schneiderman 2011). Bob does not meet the prerequisites for culpability-based desert, because it is questionable whether failing to fasten one's seatbelt is morally wrong or proportionate to any consequence relating to the health care he receives. If we remove prerequisites for culpability-based desert (e.g., by also specifying that Bob was not morally responsible for failing to fasten his seatbelt, because he was acting under hypnosis) it becomes even less plausible that he should bear any burdensome consequences. Although many would consider it too harsh to allow Celandine to suffer all the consequences of her actions, and be left to die, theorists who invoke culpability-based desert might have the resources to explain why this is so, by referring to their proportionality scale. They might argue that no wrongful act is proportionate to the outcome of being completely denied medical treatment and being left to die by the roadside. They might argue that a proportionate response would be to deprioritise Celandine's treatment in a triage situation.[16] For example, some theorists (Gold and Strous 2017) and medical professionals (Gold et al. 2021) think that after a terrorist incident victims should be treated before the injured terrorist. Alternatively, Celandine might be required to pay some financial contribution toward her health care.

So far, I have argued that (prima facie) culpability-based desert stands a better chance of helping to justify responsibility-sensitive health policies than some alternative approaches. The luck egalitarian's attempt to infer liability responsibility from outcome responsibility fails to spell out certain premises. With the help of culpability-based desert these missing premises might be filled in:

P1: X is morally responsible (to W extent) for course of conduct Y, which carries a risk of harm to X's health.

P2: Course of conduct Y is morally wrong.

P3: People who are morally responsible (to W extent) for morally wrongful course of conduct Y deserve to suffer proportionate burdensome consequence Z.

C: X is liable to undergo Z.

people can sometimes be fully morally responsible and *if* Celandine meets whatever conditions the reader believes would be needed for full moral responsibility, then it is more plausible that Celandine should take the consequences of her actions, than if she did not meet these conditions.

[16] Proportionality is functioning as what Olsaretti (2009: 167) calls "a principle of stakes", i.e., "an account of what consequences can justifiably be attached to features that are the appropriate grounds of responsibility".

Although it has more prima facie plausibility than other responsibility-sensitive approaches, ultimately this approach is problematic. The remainder of this chapter casts doubt on the idea that responsibility-sensitive health care policies based on culpability-based desert are justifiable. I challenge each of P1–P3, claiming respectively that there are reasons for doubting whether people are morally responsible for harming their own health, whether harming one's own health is morally wrong and that significant penalties for harming one's health would be a proportionate response.

3.5 Are People Morally Responsible for Harming Their Health?

A number of theorists have defended moral responsibility scepticism (e.g., Honderich 1993; Strawson 1994, 2015; Pereboom 2001; Callender 2010; Waller 2011; Vilhauer 2013; Shaw 2014; Alces 2018; Corrado 2018; Levy 2018; Focquaert 2019; Caruso 2021b). The sense of moral responsibility at issue in this literature is that which is relevant to culpability-based desert. Moral responsibility in this sense is backward-looking. Holding someone morally responsible in this sense implies that, because of a past action, they deserve a negative response, such as blame or punishment, solely because of the nature of the action and irrespective of consequentialist considerations. Moral responsibility sceptics typically endorse forward-looking senses of responsibility, such as Pereboom's (2001, 2021) account, which aims at improving an individual's future conduct by, for example, engaging in dialogue with the individual about a wrongful act that they have committed, drawing the individual's attention to what that action says about their character, providing reasons for acting differently, and so on. Moral responsibility sceptics also insist that wrongdoers (e.g., rapists or murderers) who pose a risk to others should not be allowed to roam free. However, the justification for restricting their liberty is based on forward-looking considerations, such as social protection, not on culpability-based desert. Moral responsibility sceptics would oppose penalties for self-harm based on the self-harmer's supposed culpability.[17] Some moral responsibility sceptics might endorse penalties based on utilitarian considerations, but these are beyond the scope of this chapter.

A debate has raged for millennia about whether humans possess the kind of free will that could provide a fair basis for holding people morally responsible in the backward-looking sense. One traditional challenge to the idea that humans possess this kind of free will arises from the view that all our decisions and actions

[17] "Self-harm" is used here as a shorthand for "damaging one's health", which might not constitute all-things-considered harm, e.g., if the individual values the health-damaging activity more than s/he values the aspect of her health that was damaged.

are determined by factors beyond our control, such as our biological constitution and environment. If determinism in this sense is true, factors beyond our control determine how our characters develop and which reasons we see fit to act upon. Given the influence of such factors, all our decisions (including each step in any rational deliberation leading up to them) must occur exactly as they do. These factors do not just make it more likely that we make particular decisions. They render it inevitable that we have the motivations and preferences that we have and that we reason and decide as we do. Most contemporary responsibility sceptics take a no-moral-responsibility-either-way position, according to which moral responsibility is incompatible with both determinism and indeterminism. One argument against indeterministic free will is that, rather than making us morally responsible, indeterminism would simply introduce randomness into our deliberative processes.

Another recent trend among moral responsibility sceptics is to focus on epistemic considerations (e.g., Pereboom 2001; Double 2002; Rosen 2004; Vilhauer 2009, 2013, 2023; Shaw 2014, 2021; Caruso 2018, 2021a, 2021b; Levy 2018; Corrado 2018). Although moral responsibility sceptics cannot provide 100 per cent proof that people are not morally responsible, it is hard to deny that there is significant room for doubting whether people are morally responsible. This doubt has important practical implications. It seems unfair to harm someone unless we are highly confident that harming them would be justified. This plausible notion underlies the principle—central to our criminal justice system—that people can only be punished if the case against them is proved beyond reasonable doubt. If a person's moral responsibility for performing an act is necessary to justify harming them, then according to the epistemic argument, it should be proved to a very high degree of certainty that this person was indeed morally responsible. However, the existence of the intractable debate surrounding moral responsibility suggests that there is not a sufficiently high degree of certainty that people are morally responsible, since respected experts in the field who have thought seriously about the matter fail to agree on whether people are morally responsible and on what the requirements for moral responsibility should be (Bourget and Chalmers 2014). There is a wide variety of challenges to the notion of moral responsibility (for an overview see Caruso 2018). Suffice it to say that with every hurdle new possibilities for error arise and the cumulative effect of this is to create considerable room for doubt. In Levy's (2018: 462) words, "it seems plausible that no one can confidently conclude that we know that every one of these challenges fails".

When deciding whether someone should suffer harm because they were responsible for making bad choices about their health, the severity of the harm seems relevant. The more severe the harm, the more certain we should be that the harm is justified. If the harm is very severe, we need a high degree of credence in the soundness of all elements of the argument for allowing the person to suffer it.

If the claim that they are responsible for their poor health is an element of the overall argument for harming them, we should be certain they were responsible. It is doubtful that we can be sufficiently certain, both because of the doubts about moral responsibility in general and because of the difficulties of establishing that the prerequisites for moral responsibility have been met in any particular case involving self-harm (Levy 2018; Brown et al. 2018). Harming one's health often involves going against one's strongest natural instinct (self-preservation), thwarting highly valued goals (one's own well being and flourishing), typically occurs under considerable stress and pressure, and often occurs under the influence of addiction which may have started below the age of responsibility (see, e.g. Wiss et al. 2020). Specific acts of self-harm often take place in the absence of witnesses, and damaging one's health is often associated with limited education about health issues and often coexists with or is a symptom of mental illness (see, e.g. Harwood et al. 2007). Each factor would make it difficult for medical professionals to establish with a high degree of certainty that any specific individual met the control and rationality prerequisites for moral responsibility at the specific time the individual engaged in health-harming behaviour. Given these doubts, imposing significant penalties for harming one's health seems unjustifiable, although, as I will discuss near the end of the article, the imposition of very minor burdens may be acceptable.

It might be objected that, despite these uncertainties, responsibility should be used as a "tie-breaker" in situations where one patient must lose out. For example, when there are two transplant candidates and only one suitable organ, perhaps the organ should go to the patient who seems less likely to be responsible for her illness. Indeed, the luck egalitarian might claim that failing to do so would be a "harsh" way of treating the patient who apparently was not (or was apparently less) responsible for her condition. However, responsibility would be a particularly problematic tie-breaker because, unlike other considerations, such as the likelihood that the treatment will be beneficial, attributing responsibility carries stigma. This stigma would affect patients and their loved ones when they are most vulnerable, which would be hard to justify given the uncertainties described above. Furthermore, such stigma has many counterproductive effects (Helweg-Larsen 2019)—including the impact on the doctor–patient relationship (Hansson 2018)—which could engender widespread mistrust in the health system felt by patients who fear they might be blamed, in addition to those who are actually blamed. Any value there might be in using as a tie-breaker the *possibility* that a patient *may* have been responsible for her illness could well be outweighed by the adverse consequences of doing so.[18]

[18] For discussion of other counterproductive effects of responsibility-sensitive policies besides stigma see (Levy 2018).

3.6 Is Harming One's Own Health Morally Wrong?

Three of the most influential justifications for holding people morally responsible for their own health ("pro-moral-responsibilisation arguments") have been identified and critiqued by Brown et al. (2019). They are: (1) *obligations of solidarity* to other health care system users, (2) *special relationships* (e.g., obligations to one's child or employer) and (3) *self-regarding obligations*. This section builds on Brown et al.'s critiques, addressing solidarity, special relationships and duties to oneself in turn, and provides additional reasons for doubting the soundness of these viewpoints.[19]

The *solidaristic obligations* argument appeals to the idea that public health care systems like the NHS are based on the relationship of solidarity between members of a community who share a common plight (e.g., all members face the possibility of becoming ill or injured in the future) and who decide to share the financial costs of their health risks. On this view, the creation of a publicly funded health system, based on solidarity, gives rise to reciprocal rights and duties. For example, the most vulnerable have a right not to be abandoned to their fate, but correspondingly, people who may benefit from the system have obligations to reduce costs to other health care system users by leading a healthy lifestyle. According to the pro-moral-responsibilisation version of this argument, people who fail to discharge their moral responsibility to look after themselves should be penalised. For example, smokers and the obese could be deprioritised on waiting lists for surgery.

When critiquing the solidaristic obligations argument, Brown et al. draw on two claims defended by Wilkinson (1999): first, there is little empirical evidence suggesting that those typically targeted by responsibilisation policies place a disproportionate burden on the health care system (e.g., smokers contribute to the economy via taxes on cigarettes and those who die young will not require expensive long-term care in old age); second, the logic behind penalising "lifestyle diseases" would result in an implausibly wide range of risky conduct being penalised. Furthermore, such policies may be tainted by class-based discrimination, since targeted behaviours (e.g., excessive drinking, smoking and obesity) are linked with socio-economic deprivation, while other equally risky behaviours (e.g., horse-riding, foreign travel and skiing) that are more often associated with wealthier classes are tolerated or celebrated.

I would add that, if any action can meet the control prerequisite for moral responsibility, it is likelier that this pre-condition is more frequently met when individuals engage in non-penalised activities associated with wealth than when stigmatised individuals engage in self-harming behaviours associated with poverty, since the latter are often performed under greater pressures and with the

[19] Like Levy (2018), Brown, Maslen and Savulescu claim that such arguments would only result in relatively few people being considered morally responsible for their health, as many of those who harm their health do not meet the control prerequisite for moral responsibility.

benefit of less information/education. Furthermore, if people can be held morally responsible, they are arguably more responsible or blameworthy for premeditated acts than for impulsive ones, and complex activities such as foreign travel and skiing are more likely to require careful conscious planning than purchasing cigarettes or ordering too many drinks in the pub.

My argument is not simply that responsibility-sensitive policies target the wrong behaviours. I am not proposing that health care systems should penalise skiing, horse-riding, and so on. Rather, I am raising a dilemma for responsibility-sensitive policies. If they penalise these kinds of health-damaging behaviour, these policies become impractical, intrusive and politically unviable. However, if they focus on the 'easier targets' the policies become arbitrary.

Furthermore, even if the control condition for responsibility were met when one engages in behaviour typically targeted by such policies, it is doubtful whether the solidarity argument can establish that self-harm is morally wrong as opposed to prudentially wrong, or, even if self-harm is morally wrong, it is doubtful whether it is sufficiently morally wrong to justify health-related penalties. One reason for saying that it is morally wrong to breach obligations of solidarity is because doing so is selfish: it involves according excessive weight to one's own interests. However, compared to more clear-cut cases of moral wrongdoing (e.g., intentionally harming another person for one's own financial gain), it is not straightforward to characterise people who harm themselves as selfish. The burden that self-harmers are willing to risk inflicting on themselves is typically much more severe than the burden they risk placing on any other individual to whom they owe solidary obligations. The smoker risks suffering and dying from lung cancer, whereas (even if smokers collectively cause a net loss to the NHS) the burden that any individual smoker places on any individual tax-payer will be comparatively minimal. It is plausible that it is more in accordance with the spirit of solidarity to abandon responsibilisation policies that risk stigmatising already vulnerable groups, than to impose penalties on self-harmers.

A second reason that harming one's health might be considered morally wrong is that self-harm could leave one unable to fulfil obligations stemming from *special relationships*. For example, some people who regularly drink excessively, or who damage their health through smoking or overeating, may be unable to discharge duties to their children, partner or employer. One might think that those who, through self-harm, fail to discharge such duties deserve to bear some health-related penalties. When criticising the special relationships argument, Brown et al. note that this argument is not "broadly applicable" as not everyone will be in a relevant relationship and obligations arising from such relationships would typically only require quite low levels of health to be discharged.

I would add that the class of cases in which it could be established that the patient was sufficiently morally responsible for a sufficiently wrongful act would be further narrowed due to the complex nature of both self-harm and special

relationships. Loving parents who self-harm are often acting irrationally by thwarting the goals they value most: their own well being and that of their child. This raises questions about whether the rationality prerequisite for moral responsibility has been met (for a related argument, see Nir Eyal's chapter in this volume). In other cases, problems within the relationship, for example, domestic abuse or unfair and stressful work conditions, are the cause of the self-harm. In such cases, further penalising the abused spouse or employee to vindicate the supposed rights of the abusive spouse or employer would be unethical. The costs in time and resources it would take for the authorities to get clear about such issues could well be greater than any benefit the health system could derive from penalising the minority of self-harmers who could be proved to have been fully responsible for wrongfully breaching obligations within special relationships, even assuming that such proof is ever possible.

Cases where the breach of duty has severe consequences (e.g., allowing a child to come to harm) are already covered by criminal laws.[20] To impose additional health-related sanctions unfairly punishes someone twice. This consideration is recognised by the principle that prisoners should receive a level of health care equivalent to that provided to other members of the community (National Prison Healthcare Board 2018).[21] Moreover, imposing health-related sanctions for breaching obligations stemming from special relationships would likely penalise the very people whose rights the policy aims to protect, that is, those with whom the self-harmer is in a special relationship, such as family members. Family members are likely to be unpaid carers for the self-harming person. Regardless of whether the health system deems the patient less deserving of care, family members are unlikely to take the luck egalitarian line of "you made your bed so you must lie in it" and abandon their relative. So, for example, if a smoker were sanctioned by being deprioritised on a waiting list for surgery, this would likely result in a longer period in which the patient would face disability and discomfort which would intensify the caring burden placed on family members. This policy could also interfere with the patient's ability to care for their children whose care might then become an additional burden on other relatives. Thus, the policy would likely penalise the very groups whose interests supposedly justified it. Furthermore, such a policy might involve indirect discrimination[22] against women, as female relatives comprise the majority of unpaid carers (Brimblecombe et al. 2017) and such a policy could therefore have a disproportionate impact on them.[23]

The third argument for holding people morally responsible for their own health is based on the Kantian idea of *duties to oneself*. For example, on this view, a person might owe it to herself to nurture her talents. A person who smokes,

[20] *R v Gibbins and Proctor* (1918) 13 Cr App R 134.
[21] See also: United Nations General Assembly (1991: Principle 9).
[22] Equality Act 2010, s.19.
[23] For discussion of discrimination against unpaid carers see Tribe and Lane (2017).

overeats or drinks excessively might develop health conditions that interfere with her ability to nurture her talent for athletic, intellectual or artistic pursuits. In reply to this argument, Brown et al. note that Kantian claims about self-regarding duties are contentious and that even if individuals have self-regarding moral obligations, it is unclear whether these can be "legitimately enforced by the state" (Brown et al. 2019). I would add that penalising failures to fulfil duties to oneself would arguably be self-contradictory and incoherent. If X has self-regarding duties, it seems plausible that the existence of such duties must rely at least partly on the idea that X's flourishing is a good thing. It thus makes little sense to respond to X's failure to promote their own flourishing by taking measures that further interfere with their flourishing. One might reply that threatening to penalise someone for self-harm could be an effective deterrent. However, there is a lack of empirical evidence for that claim; in fact, stigma and threats seem to make self-harm more likely (Hansson 2018; Helweg-Larsen et al. 2019; Turner et al. 2020). Furthermore, arguments based on deterrence seem to appeal to consequentialist reasoning that go against the Kantian rationale behind self-regarding obligations.

3.7 Are Health-Related Sanctions a Proportionate Response to Harming One's Own Health?

The most highly-developed theoretical discussion of proportionate sanctions is in the literature on criminal punishment. A similar concept of proportionality is needed to justify penalising self-harm with reference to culpability-based desert. As argued above, without a proportionality constraint such policies would lead to health outcomes that intuitively seem too harsh in certain cases. Indeed, the current section argues that there are reasons for doubting that significant penalties for harming one's health could ever be a proportionate response (although very minor burdens might be less open to challenge).

In the context of imposing sanctions as a deserved response to morally culpable conduct, "proportionality" is a fittingness relationship between the sanction and the severity of the conduct, such that more serious misconduct attracts harsher penalties. Some think that the *kind* of penalty imposed should also be appropriate to the *kind* of wrongdoing, as reflected in the biblical idea of "an eye for an eye and a tooth for a tooth". In modern societies which reject physical punishments it is seldom deemed ethically acceptable to achieve a very close match in terms of *kinds* of misconduct and penalty.[24] An acceptable example might be requiring someone who had vandalised property to do community service that involved repairing

[24] Health-related sanctions are arguably a kind of physical punishment by omission. This raises concerns about degrading treatment (article 3 of the European Convention on Human Rights) which are beyond the scope of the current chapter. See Boruckiee (2019).

property. However, a penal system in which fines and imprisonment are typical punishments clearly prioritises proportionality of severity over proportionality of kind. Indeed, prioritising proportionality of severity seems intuitively plausible, as the injustice of disproportionate punishments such as life imprisonment for a minor assault seems to lie primarily in the mismatch of severity between crime and punishment, rather than a mismatch in "kind": assaulting the offender would not be an improvement.

Indeed, proportionality of kinds can sometimes be positively misleading, opening the door to disproportionality of severity, and may be a source of confusion in the responsibility for health debate. At first glance, it may seem fitting to respond to a person who harms their health by imposing a health-related sanction, such as deprioritising their treatment. However, such a response may be disproportionate in terms of severity and also may fail to fit with the underlying justification for holding the individual morally responsible for harming her health. Regarding severity: health-related sanctions such as delaying surgery or asking people to contribute to the costs of health care are significantly burdensome. Although they are less extreme than outright treatment denial, they come with substantial costs in terms of discomfort, reduced independence, stress and stigma at a time when patients are very vulnerable and when the policy's adverse impact on mental health might affect the physical outcomes. These sanctions seem disproportionate given the reasons discussed in the previous section for doubting that self-harm is (seriously) morally wrong and for doubting that any particular individual is (fully) morally responsible for self-harm.

Regarding the fit between the sanction and the underlying justification for holding the individual morally responsible, consider the solidarity justification, arguably the strongest of the three justifications for penalising self-harm discussed above (Davies and Savulescu 2019).[25] The solidarity argument is based on the idea that people who self-harm wrong their fellow citizens by placing an undue burden on the publicly funded health care system. However, it is puzzling why the scope of one's responsibility to one's fellow citizens should be delimited in this way (beyond a misleading *semantic* match, with "health" featuring as an element of the supposed wrongdoing and an as element of the sanction). There does not seem to be a strong *moral* reason for focusing narrowly on a putative solidary obligation to safeguard one's own health and to ignore how one conducts oneself as a citizen more generally. Imagine Carl, who helps to promote the health of others, for

[25] It is not entirely clear how (or whether) Brown et al.'s critique of this argument is compatible with Davies and Savulescu's (2019) "golden opportunities" argument, which defends imposing (albeit low-level) sanctions on those who breach solidary obligations in the context of health care. It is possible that these two articles can be reconciled with each other, on the basis that the former article opposes policies that in "a wide and unrestricted way" penalise breaches of solidary obligations, whereas the latter article defends restricted, low-level penalties for a narrowly defined set of patients who reject "golden opportunities" to preserve their health.

example, through cooking healthy meals for his children and educating them about healthy habits, doing voluntary work that promotes health in his community, relieving his spouse of other domestic duties so she has more time to exercise, and working as a conscientious doctor. However, Carl (perhaps due to the strain of his altruistic activities) has certain habits that damage his own health. Now, compare this individual to Damien, who does not have a social conscience, leads a selfish lifestyle, does as little as he can legally get away with to promote the health of his family or anyone else, and profits from selling or manufacturing unhealthy products, but takes perfect care of his own health. The narrow focus on damaging one's own health adopted by the solidarity approach discussed above would impose sanctions on Carl but not Damien. However, it is unclear why that is just. I am not suggesting that the state should do a wholesale investigation of each citizen's civic virtue and allocate health resources accordingly. But the impracticality and intrusiveness of doing so does not justify focusing narrowly on people who damage their own health. Such a narrow focus is unjustifiable because it would be arbitrary (as well as probably being almost as impractical and intrusive as investigating other aspects of civic virtue).[26] The arbitrariness of the narrow focus on damaging one's own health comes on top of the arbitrariness of policies that only penalise those who take risks with their health and are *unlucky enough to get sick*, while others who took the same risk have a stronger constitution and so face no penalty.

It might be objected that it would not be arbitrary for the health system to focus specifically on penalising people who damage their own health, because it is reasonable for there to be a division of labour in the way society administers sanctions, with different institutions/sectors (e.g., the civil courts and the criminal courts, etc.) dealing with different types of breaches of duty. In reply, first, it seems arbitrary for the health system to penalise damaging/neglecting one's own health but not to penalise damaging/neglecting other people's health, given that the health system seeks to promote health in general. Second, it seems arbitrary to endow a social institution with the power to impose sanctions for damaging one's own health given the doubts discussed in this chapter about whether such conduct is morally wrong, and given that there are many other types of conduct which seem morally much worse (e.g., Damien's conduct in the above example) which no institution has the power to penalise. Third, it seems arbitrary for one social institution (the health system) to allocate burdens as severe as illness or death based on a justification (the supposed immorality of damaging one's health) which

[26] Davies and Savulescu's (2019) and De Marco, Savulescu and Douglas's (2021) account of situations in which patients could be presented with a "golden opportunity" to change their lifestyles provide the most promising attempts at describing ways of taking responsibility into account while seeking to minimise intrusiveness, unfairness and practical difficulties. However, as they acknowledge, their approach still faces challenges. One difficulty is that in order to minimise unfairness, medical professionals must consider if the patient had a good justification or excuse for harming their health. However, these justifications and excuses seem both easy for patients who understand the system to fake and difficult for the most vulnerable patients to articulate.

is so debatable, whereas other social institutions (e.g., civil and criminal courts) impose sanctions that are much less severe (e.g., fines) for behaviour whose wrongfulness is much less disputable (rightly deeming more severe sanctions to be disproportionate).

This section now responds to some potential counterarguments against my claim that responsibility-sensitive health care policies are likely to create disproportionate outcomes. Albertsen and Nielsen (2020) have argued that it is incoherent to rely on the concept of "disproportionality" when criticising responsibility-sensitive health care policies, because one cannot invoke disproportionality without implicitly endorsing the very responsibility-sensitive approach that one was attempting to criticise. They also claim that a proportionality principle is not capable of ruling out harsh cases, since to say that it is disproportionate to attach certain burdens to certain instances of imprudence "must imply that responsibility-sensitivity in *some proportion* is appropriate for distributive justice. It also implies that we can imagine situations involving extremely bad consequences for the imprudent as a result of a proportionally bad exercise of responsibility which therefore does not disturb justice from the perspective of proportionality" (2020: 6).

However, neither of these supposed implications follows from the claim that it would be disproportionate to penalise people for poor health decisions.

In reply to their first claim (the incoherence claim): one can criticise a theory by appealing to a principle that is part of that theory without endorsing the principle or theory. The critic might simply be noting that the theory is internally inconsistent because its implications do not fit with its own principles. Thus, a critic of a theory that purports to attach proportionate consequences to responsible choices can coherently argue that certain consequences would be "disproportionate", without accepting the validity of this responsibility-sensitive theory and without accepting the proportionality principle. For example, if a particular retributive theory attempted to justify punishing shoplifting with death, critics could argue that this would be "disproportionate" by retributivism's own lights. Critics could coherently make this point without endorsing either retributivism or the retributive concept of proportionality. Similarly, critics might argue that imposing severe health consequences for imprudence would violate the proportionality principle that proponents of responsibility-sensitive health policies endorse without themselves endorsing responsibility-sensitivity or this proportionality principle. Alternatively, the critic might endorse proportionality, but interpret this concept in a way that does not depend on responsibility. For instance, some writers endorse the idea that a burden placed on an individual must be proportionate to the severity of an ongoing threat posed by that individual, regardless of the individual's responsibility status (e.g., Honderich 1993; Pereboom 2001; Caruso 2018; Shaw 2019). A carrier of an infectious disease might not be responsible for having that disease, but it might nevertheless be proportionate to quarantine that

individual if the disease were sufficiently dangerous and infectious and if the quarantine were not too burdensome, and disproportionate to quarantine them if those conditions were not met. Finally, a critic might endorse responsibility-sensitive policies in certain contexts (e.g., criminal punishment), while coherently arguing that imposing burdens on people for damaging their own health is always disproportionate because such conduct is not wrongful enough to warrant state interference.

In reply to Albertsen and Nielsen's second claim, that proportionality principles cannot rule out harsh cases: it is unclear why they assume that a scale of proportionate burdens can have no upper limit. Albertsen and Nielsen (2020: 6) state that "we can imagine situations involving extremely bad consequences for the imprudent as a result of a proportionally bad exercise of responsibility which therefore does not disturb justice from the perspective of proportionality". It is puzzling why they do not consider the possibility that a proponent of proportionality might deem it unjust (because disproportionate) to impose "extremely bad consequences" on someone for mere imprudence no matter the extent of the imprudence. Analogously, a penal system might have a punishment scale that increased in severity in proportion to the severity of the crime until the maximum punishment allowable in that jurisdiction was reached (say life imprisonment). The penal system might presuppose that there is nothing a person can do in response to which a punishment more severe than life imprisonment would be proportionate. So too, the health care system could presuppose that there is no form of imprudence in response to which imposing/allowing "extremely bad consequences" would be proportionate. Of course, it has frequently been argued (often by critics of retributivism) that there is no principled way of devising a proportionality scale. Reasonable people can disagree about which kind of burden is a proportionate response to which kind of conduct and there is no obvious way to settle this disagreement. However, most people can agree about certain clear cases of disproportionality (e.g., the death penalty for theft, or being left to die because one didn't buckle one's seatbelt) and so the concept of proportionality (or gross disproportionality) might still be useful for ruling out these harsh cases. Furthermore, to return to the coherence point discussed in the previous paragraph, it would be coherent to object to responsibility-sensitive policies because (1) these policies rely on the notion of proportionality, (2) these policies imply that people should bear burdens that would be disproportionate *if* the concept of disproportionality were sound and (3) the concept of disproportionality is flawed because there is no principled way to devise a justifiable proportionality scale.

I have suggested that it would be disproportionate to impose significant penalties for harming one's health, such as deprioritising treatment for smokers and the obese or making such individuals contribute to the costs of their treatment. However, minor burdens might be less open to challenge. For example, taxes on cigarettes and minimum alcohol pricing may withstand the objections raised in this chapter. They are likely to be less stigmatic and emotionally damaging than

other policies that impose burdens on those engaging in unhealthy behaviour, such as policies that cite a person's past irresponsible conduct as a reason for delaying surgery. Taxes and minimum pricing affect individuals at the point when they take the risk, rather than arbitrarily penalising only those risk-takers who are unlucky enough to get sick. They are not penalties that are only imposed when patients are at their most vulnerable. Taxing and minimum pricing also affect companies who sell unhealthy products, not just consumers, unlike policies which disproportionately impact socially disadvantaged groups and leave companies who profit from selling unhealthy goods untouched. Arguably, taxation and minimum pricing engage with a more forward-looking sense of responsibility, which is less vulnerable to the challenge from moral responsibility sceptics than the backward-looking moral responsibility that features in other potential policies. Because taxing and minimum pricing have their impact at the time when the consumer is deciding whether to take the risk, they have the potential to encourage the consumer to engage in more responsible future behaviour. In contrast, sanctions such as treatment deprioritisation that penalise the individual after the self-harm has already happened are more dependent on the idea of backward-looking responsibility (and any deterrent value they may have may be undermined by the counterproductive effects of stigma).

3.8 Conclusion

This chapter discussed some of the difficulties faced by proponents of responsibility-sensitive approaches in health care, highlighting the implications of moral responsibility scepticism. When discussing punishment, moral responsibility sceptics have often invoked an epistemic argument, maintaining that there are at least serious doubts about whether people are morally responsible (in the sense required for retributive punishment) and that, in view of this uncertainty, retributive punishment is unjust, given the serious harm it inflicts on offenders. This type of reasoning also implies that we should not take patients' responsibility for their poor health into account when deciding whether to give these patients treatments for serious illnesses and that the health system should not impose significant penalties on individuals for harming their own health. However, minor burdens, such as taxes on cigarettes and minimum alcohol pricing might be less open to challenge.

References

Abad, D. (2011), 'Desert, responsibility and luck egalitarianism', in N. Vincent, I. van de Poel and J. van den Hoven (eds), *Moral Responsibility* (Springer): 121–40.

Albertsen, A., and Knight, C. (2015), 'A framework for luck egalitarianism in health and healthcare', in *Journal of Medical Ethics* 41/2: 165–9.

Albertsen, A., and Nielsen, L. (2020), 'What is the point of the harshness objection?' in *Utilitas* 32/4: 427–43.

Alces, P. (2018), *The Moral Conflict of Law and Neuroscience* (University of Chicago Press).

Anderson, E. (1999), 'What is the point of equality?', in *Ethics* 109/2: 287–337.

Arneson, R. (2000), 'Luck egalitarianism and prioritarianism', in *Ethics* 110/2: 339–49.

Borucki, C. (2019), 'When healthcare goes up in tobacco smoke: A selective healthcare system from a (European) human rights perspective', in *Utrecht Law Review* 15/3: 6–26.

Bourget, D., and Chalmers, D. (2014), 'What do philosophers believe?', in *Philosophical Studies* 170/3: 465–500.

Brimblecombe, N., Pickard, L., King, D. and Knapp, M. (2017), 'Perceptions of unmet needs for community social care services in England: A comparison of working carers and the people they care for', in *Health and Social Care in the Community* 25/2: 435–46.

Brouwer, H., and Mulligan, T. (2019), 'Why not be a desertist?' in *Philosophical Studies* 176/9: 2271–88.

Brown, R., Maslen, H. and Savulescu J. (2018), 'Responsibility, prudence and health promotion', in *Journal of Public Health* 41/3: 561–5.

Brown, R., Maslen, H. and Savulescu, J. (2019), 'Against moral responsibilisation of health: Prudential responsbility and health promotion', in *Public Health Ethics* 12/2: 114–29.

Callender, J. (2010), *Free Will and Responsibility: A Guide for Practitioners* (Oxford University Press).

Caruso, G. (2018), 'Skepticism about moral responsibility', in E. Zalta (ed.), *The Stanford Encyclopaedia of Philosophy*. Available at: https://plato.stanford.edu/entries/skepticism-moral-responsibility/ [Accessed 18/08/22].

Caruso, G. (2021a), 'Retributivism, free will skepticism and the public health-quarantine model: Replies to Corrado, Kennedy, Sifferd, Walen, Pereboom and Shaw', in *Journal of Legal Philosophy* 46/2: 161–215.

Caruso, G. (2021b), *Rejecting Retributivism: Free Will, Punishment and Criminal Justice* (Cambridge University Press).

Corrado, M. (2018), 'Free will, punishment and the burden of proof', in *Criminal Justice Ethics* 37/1: 55–71.

Cushman, F. (2008), 'Crime and punishment: Differential reliance on causal and intentional information for different classes of moral judgment', in *Cognition* 108: 353–80.

Davies, B. (2022), 'Responsibility and the recursion problem', in *Ratio* 35/2: 112–22.

Davies, B., and Savulescu, J. (2019), 'Solidarity and responsibility in health care', in *Public Health Ethics* 12/2: 133–44.

De Marco, G., Douglas, T. and Savulescu, J. (2021), 'Healthcare, responsibility and golden opportunities', in *Ethical Theory and Moral Practice* 24: 817–31.

Double, R. (2002), 'The moral hardness of libertarianism', in *Philo* 5/2: 226–34.

Duff, R. (2001), *Punishment, Communication and Community* (Oxford University Press).

East and North Hertfordshire Clinical Commissioning Group (2017), 'East and North Hertfordshire NHS service changes: Decisions announced'. Available at: https://www.enhertsccg.nhs.uk/news/201710/east-and-north-hertfordshire-nhs-service-changes-%E2%80%93-decisions-announced [Accessed 7/04/22].

Equality Act (2010), (c.15). Available at https://www.legislation.gov.uk/ukpga/2010/15/contents [Accessed 30/08/22].

Eyal, N. (2006), 'Egalitarian justice and innocent choice', in *Journal of Ethics and Social Philosophy* 2/1: 1–18.

Feinberg, J. (1970), *Rights and Reason* (Springer).

Feldman, F. and Skow, B. (2020), 'Desert', in E. Zalta (ed.), *The Stanford Encyclopaedia of Philosophy*. Available at: https://plato.stanford.edu/entries/vagueness/ [Accessed 18/08/22].

Fischer, J. and Ravizza, M. (1998), *Responsibility and Control: A Theory of Moral Responsibility* (Cambridge University Press).

Focquaert, F. (2019), 'Free will skepticism and criminal punishment: A preliminary ethical analysis', in E. Shaw, D. Pereboom and G. Caruso (eds), *Free Will Skepticism in Law and Society: Challenging Retributive Justice* (Cambridge University Press): 207–36.

Gold, A., Greenberg, B., Strous, R. and Asman, O. (2021), 'When do caregivers ignore the veil of ignorance? An empirical study on medical triage decision–making', in *Medicine, Health Care and Philosophy* 24/2: 213–25.

Gold, A., and Strous, R. (2017), 'Second thoughts about who is first: The medical triage of violent perpetrators and their victims', in *Journal of Medical Ethics* 43: 293–300.

Hansson, S. (2018), 'The ethics of making patients responsible', in *Cambridge Quarterly of Healthcare Ethics* 27: 87–92.

Hart, H. (1968), *Punishment and Responsibility* (Oxford University Press).

Harwood, G., Salsberry, P., Ferketich, A. and Wewers, M. (2007), 'Cigarette smoking, socioeconomic status, and psychosocial factors: Examining a conceptual framework', in *Public Health Nursing* 24/4: 361–71.

Health and Social Care Act (2012), (c. 7). Available at https://www.legislation.gov.uk/ukpga/2012/7/contents/enacted [Accessed 30/08/22].

Health and Social Care Act 2022 (c. 31). Available at https://www.legislation.gov.uk/ukpga/2022/31/contents/enacted [Accessed 30/08/22].

Helweg-Larsen, M., Sorgen, L. and Pisinger, C. (2019), 'Does it help smokers if we stigmatize them? A test of the stigma-induced identity threat model among US and Danish smokers', in *Social Cognition* 37/3: 294–313.

Honderich, T. (1993), *How Free Are You?* (Oxford University Press).

Kane, R. (1998), *The Significance of Free Will* (Oxford University Press).

Lasersohn, P. (1999), 'Pragmatic halos', in *Language* 75/3: 522–51.

Levy, N. (2018), 'Responsibility as an obstacle to good policy: The case of lifestyle related disease', in *Journal of Bioethical Inquiry* 15/4: 459–68.

Lippert-Rasmussen, K. (2016), *Luck Egalitarianism* (Bloomsbury).

Moore, M. (1997), *Placing Blame* (Oxford University Press).

National Prison Healthcare Board. (2018), 'Principle of Equivalence of Care for Prison Healthcare in England'. Available at https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/837882/NPHB_Equivalence_of_Care_principle.pdf [Accessed 24 August 2022].

Olsaretti, S. (2009), 'Responsibility and the consequences of choice', in *Proceedings of the Aristotelian Society* 109/1: 165–88.

Parikh, R. (1994), 'Vagueness and utility: The semantics of common nouns', in *Linguistics and Philosophy* 17: 521–35.

Pereboom, D. (2001), *Living without Free Will* (Cambridge University Press).

Pereboom, D. (2021), 'Undivided forward-looking moral responsibility', in *The Monist* 104: 484–97.

Pillutla, V., Maslen, H. and Savulescu, J. (2018), 'Rationing elective surgery for smokers and obese patients: Responsibility or prognosis?', in *BMC Medical Ethics* 19/1: 1–10.

R v Gibbins and Proctor (1918), 13 Cr App R 134.

Ristroph, A. (2006), 'Desert, democracy, and sentencing reform', in *Journal of Criminal Law and Criminology* 96/4: 1293–352.

Rosen, G. (2004), 'Skepticism about moral responsibility', in *Philosophical Perspectives* 18: 295–313.

Royal College of Surgeons (2016), 'Smokers and overweight patients: Soft targets for NHS savings'. Available at https://www.rcseng.ac.uk/-/media/files/rcs/library-and-publications/non-journal-publications/smokers-and-overweight-patients–soft-targets-for-nhs-savings.pdf [Accessed 23/06/22].

Schneiderman, L. (2011), 'Response to open peer commentaries on "Rationing just medical care"', in *The American Journal of Bioethics* 11/10: 1–3.

Schneiderman, L., and Jecker, N. (1996), 'Should a criminal receive a heart transplant? Medical justice vs societal justice', in *Theoretical Medicine* 17: 33–44.

Segall, S. (2015), 'What's so egalitarian about luck egalitarianism?' in *Ratio*, 28/3: 349–68.

Shaw, E. (2014), *Free Will, Punishment and Criminal Responsibility* (PhD thesis, University of Edinburgh). Available at: https://era.ed.ac.uk/bitstream/handle/1842/9590/Shaw2014.pdf?sequence=2&isAllowed=y [Accessed 24/08/22].

Shaw, E. (2019), 'Justice without moral responsibility?' in *Journal of Information Ethics* 28/1: 95–114.

Shaw, E. (2021), 'The epistemic argument against retributivism', in *Journal of Legal Philosophy* 46/2: 155–60.

Sorensen, R. (2022), 'Vagueness', in E. Zalta (ed.), *The Stanford Encyclopaedia of Philosophy*. Available at: https://plato.stanford.edu/entries/vagueness/. [Accessed 18/08/22].

Strawson, G. (1994), 'The impossibility of moral responsibility', in *Philosophical Studies* 75/1: 5–24.

Strawson, G. (2015), 'The impossibility of ultimate responsibility?' in J. Dancy and C. Sandis (eds), *Philosophy of Action: An Anthology* (Wiley), 373–81.

Tribe, R., and Lane, P. (2017), 'Anti-discriminatory practice: Caring for carers of older adults with mental health dilemmas', in P. Lane and R. Tribe (eds) *Anti-discriminatory Practice in Mental Health Care for Older People* (Jessica Kingsley) 147–74.

Turner, M., Ford, L., Somerville, V., Javellana, D., Day, K. and Lapinski, M. (2020), 'The use of stigmatizing messaging in anti-obesity communications campaigns: Quantification of obesity stigmatization', in *Communication Reports* 33/3: 107–20.

United Nations General Assembly (1991), *Basic Principles for the Treatment of Prisoners*, UN Doc A/RES/45/111 (28 March 1991). Available at https://www.ohchr.org/sites/default/files/basicprinciples.pdf [Accessed 30/08/22].

Vilhauer, B. (2009), 'Free will and reasonable doubt', in *American Philosophical Quarterly* 46/2: 131–40.

Vilhauer, B. (2013), 'Persons, punishment, and free will skepticism', in *Philosophical Studies* 162/2: 143–63.

Vilhauer, B. (2023), 'Free will skepticism and criminals as ends in themselves', in M. Altman (ed.), *Palgrave Handbook on the Philosophy of Punishment* (Palgrave Macmillan).535-556.

Vincent, N. (2009), 'What do you mean I should take responsibility for my own ill health', in *Journal of Applied Ethics and Philosophy* 1: 39–51.

Walen, A. (2020), 'Retributive justice', in E. Zalta (ed.), *The Stanford Encyclopaedia of Philosophy*. Available at: https://plato.stanford.edu/entries/justice-retributive/ [Accessed 18/08/22].

Waller, B. (2011), *Against Moral Responsibility* (MIT).

West, J., Chao, S., Kelley, S., Schwartz, J., Bertsch, D. and Marsh, J. (2003), 'Organ allocation: A case for not transplanting the violent criminal', in *Seminars in Dialysis* 16/5: 362–4.

Wilkinson S. (1999), 'Smokers' rights to health care: Why the "restoration argument" is a moralising wolf in a liberal sheep's clothing', in *Journal of Applied Philosophy* 16: 255–69.

Wiss, D., Avena, N. and Gold, M. (2020), 'Food addiction and psychosocial adversity: Biological embedding, contextual factors, and public health implications', *Nutrients* 12/11: 3521. Available at https://doi.org/10.3390/nu12113521 [Accessed 24/08/22].