

Responsibility for Ill-Health and Lifestyle

Drilling Down into the Details

Neil Levy

7.1 Introduction

Whether agents are morally responsible for their need for scarce resources is a difficult and fraught issue. It's not simply because the question whether agents are morally responsible at all is itself contentious (though it is; moral responsibility scepticism is an increasingly influential view: Caruso 2021; Pereboom 2014; Shaw et al. 2019). It is also because showing that agents are sometimes morally responsible for some of their actions—even for some of the actions that help to cause their ill-health—isn't anywhere near enough to show that they are morally responsible for their need for scarce resources. In this paper, I aim to explore some unappreciated difficulties for the attribution of moral responsibility for needs that arise from the fact that in typical cases, ill-health arises from lifestyle: not, that is, from one bad decision, but from a long-term pattern of actions.

My aim in exploring this issue is twofold. First, I hope to build on Brown and Savulescu's (2019) programmatic exploration of what they call the diachronic condition on moral responsibility for ill-health. I will show that the diachronic condition fractionates in multiple ways, depending on how ill-health is caused as well as on the theory of moral responsibility at issue. Second, I aim to show that it is much harder to satisfy the diachronic condition on moral responsibility for ill-health than is widely assumed. I will argue that we usually cannot be confident that a particular agent (or a class of agents: say alcoholics) is responsible for their ill-health, and that—I will suggest—should make us hesitant to ascribe responsibility to them.

However, there is a feature of responsibility-ascriptions in the health care context that distinguishes it from some other contexts in which we ascribe responsibility, such as the criminal justice system. This feature might render the foregoing, if not quite moot, at least very much muted. In the criminal justice

context, we are often deciding whether *anyone* is to suffer a burden.¹ Will someone go to jail or be publicly shamed, for instance? What is at stake is whether someone will be punished. This fact entails that we face a heavy justificatory burden: only if we have shouldered that burden may we proceed to holding the person responsible. But matters are quite different in the context at issue here. When we are dealing with the allocation of scarce resources, the question is not *whether* a burden will be borne, but *who* will bear a burden. Will it be patient A or patient B? Since the question is not whether but who, we may think we need much weaker grounds for settling the question. We needn't meet the threshold the criminal law sets itself in order to be justified in holding agents responsible in contexts like these.

We face, as it were, a different burden of proof in virtue of this difference in the imposition of burdens. In the criminal law, the prosecution has to prove guilt beyond reasonable doubt. The set of conditions just mentioned may be taken to show that we will typically struggle to reach this threshold: there will usually be reasonable doubt whether a particular individual or class of individuals satisfies these conditions. But surely we will often be able to show that *on the balance of probabilities*, the individual or group is responsible? Isn't that enough to show that we can use responsibility as a tie-breaker? Suppose we must allocate a scarce resource (say a kidney) to either patient A or patient B, each of whom has roughly the same medical need for it. Evidence that A is more probably than not responsible for needing the kidney, whereas B is not, might be enough to tip our decision B's way.

It is to this issue I will turn in the final section of this paper. Having shown that we rarely can shoulder a reasonable doubt burden of proof with regard to needs for scarce resources, I will assess the prospects for a balance of probabilities standard. Can we establish that even if we cannot be confident that an agent deserves to bear a burden, we may nevertheless be confident that she deserves to bear it *more* than does some other agent? I will argue that even if we have the right to such confidence, this may not be enough for responsibility to serve as a tie-breaker between agents in the context of the allocation of scarce resources.

7.2 Responsibility and the Allocation of Resources

It is a familiar fact that agents may be causally responsible for their own need for scarce health-related resources. A large and growing proportion of the global

¹ Ben Davies has pointed out that this is an oversimplification. There are at least two contexts in which the criminal justice system might be faced with deciding who a burden falls on, rather than whether a burden falls on anyone. One concerns harms to the victims of past wrongdoing: such victims or their families may experience psychological harms stemming from the acquittal of someone they believe (or even know) to be guilty. Another context concerns offenders with a high chance of recidivism: if they are not burdened now, a future victim may be instead. Our unwillingness to relax our standards for a justified finding of responsibility in the criminal law in these contexts may stem in part from a perceived difference in the magnitude and the certainties of the harms that acquittal risks.

burden of disease is lifestyle related (Yoon et al. 2014). For example, it has been estimated that nearly 40 per cent of all cancers diagnosed in the UK could be prevented by changes in lifestyle (Davis 2018). Smoking, which significantly raises the risk for lung cancer, as well as multiple other cancers, is the single biggest lifestyle contributor to cancer risk, but obesity, excessive alcohol consumption, overconsumption of certain foods (like processed meats) and physical inactivity are all risk factors for other cancers. Most of these same factors are also risk factors for other life-threatening conditions like stroke and cardiopulmonary disease. In other words, the big killers—those that account for most deaths in developed countries—are all very sensitive to lifestyle factors, and many deaths are premature as a consequence.

In addition, lifestyle contributes significantly to a range of other problems. Unprotected sex is obviously a risk factor for sexually transmitted disease. Alcohol consumption may lead to kidney disease, and other drugs can lead to a range of health problems ranging from minor to life-threatening. Lifestyle also contributes to mental health problems, in addition to addiction. There is evidence, for instance, that psychosis may be linked to use of cannabis (Di Forti et al. 2014; Volkow et al. 2016).

Health systems therefore very often see patients who need scarce resources as a consequence of their lifestyle. These patients may be chronically ill (for example, with emphysema) and require ongoing care, in their own homes or in hospitals. They may suffer from liver disease and be in need of a transplant. People struggling with addictions and other mental health problems may need services outside the clinical setting, narrowly construed; services ranging from housing through to counselling and job training. Decent societies provide these kinds of care and services to those who need them when they cannot pay for them.

But we may wish to distinguish some potential recipients of these services and care from others. All these resources are, even in the most generous and caring of societies, limited. They cost money, and money is a limited resource. Whatever resources we expend on health care for these patients are not available to others (sick children, say) and the money spent is not available for other ends (funding research; the arts; pensions, and so on). All expenditures have opportunity costs, where the opportunity cost of an expenditure is the best alternatives use of the expenditure. Other resources that are not, or not directly, financial are even more obviously limited: physician's time and attention, for example. Of course many non-financial resources could be expanded given extra financial resources: we could, for example, pay for more hospital beds or for the training of more physicians. But some resources are in inelastic supply: while we could increase supply of kidneys or livers (the latter through living donation) through financial incentives, supply of other organs is harder to increase by increasing expenditure.²

² Of course, many donatable organs are buried or cremated, either because the deceased person or their family is opposed to donation or because they simply failed to express a preference. Opposition of

Because health resources are limited, health care providers are inevitably faced with difficult resource allocation decisions. Every dollar spent on one patient's care might have been spent on another. The subsidy of a class of drug entails that others are not subsidized (or that bed numbers are not increased, and so on). Of course, the need to wrestle with opportunity costs is absolutely pervasive. We all face it, in all domains: going to the cinema entails not going to a restaurant; reading a book entails not reading another book; talking to friends entails not meditating, and so on. The domain of health is not special in entailing opportunity costs nor in entailing difficult decisions (many people buy houses; settling on a particular house is for most extremely stressful, precisely because of the enormous opportunity costs of the biggest expenditure most people will ever make). But in health care, the stakes are literally life and death; which is to say that the opportunity costs of every expenditure are the deaths of other individuals. This is true not only where it is obvious (emergency rooms might have to engage in triage, and as a consequence patients may die who would otherwise have lived) but *everywhere*: while investment in physician numbers or drugs or for that matter public information campaigns have a diminishing marginal utility, it will always be true that had we spent more money *in this way* rather than *that*, people who died might have lived.

It is widely held (even by people who reject consequentialism as the right approach to morality) that the best way to make such life-and-death decisions is via some kind of utilitarian calculus (it is such an approach, for instance, that lies behind the widespread use of QALYs in health allocation decisions and policy; McKie et al. 1998; Neumann et al. 2014). On this view, we should allocate scarce resources to get the best bang for our buck, with "bang" measured in QALYs or via some other mechanism. While all decisions entail opportunity costs, some decisions are better than others. We should choose to expend resources such that opportunity costs are minimized.

While this kind of approach to resource allocation decisions obviously has much to recommend it, even many of its advocates believe that it should be modified. It is at this point that responsibility enters the picture. While we ought to allocate resources to get the best bang for our buck, other things equal, things are not equal when some, but not all, the potential beneficiaries of these resources are responsible for their need for them. On the influential luck egalitarian view, for instance, individuals deserve compensation for the consequences of their bad (brute) luck, when it makes them worse off than others, but not for the consequences of their own decisions (e.g., Cohen 2011). Translated into the context of health care, luck egalitarianism entails that agents who are responsible for their own needs should be given a lower priority than those who are unlucky in needing scarce resources (Segall 2010; in this volume, however, Eyal argues that

the first kind and inertia of the second would both fall given significant financial incentives. It's unlikely that such a measure would eliminate the need to make hard choices between potential recipients of organs, due to the inherent difficulties in transporting organs very far.

luck egalitarianism need not support this conclusion). One need not be a luck egalitarian to believe that health care should be sensitive to responsibility: the view is widely held and may be intuitive for many, and it is enshrined in health policy (see Brown et al. 2019 for discussion).

Importantly, one need not believe that those who are responsible for their ill-health deserve to suffer to find it intuitive that resource allocation should be sensitive to considerations of responsibility. Retributivists believe that the guilty deserve to suffer; the responsibility might take the analogous line in the domain of health care (retributivism about lifestyle is not a mere thought experiment: think of how medieval Christianity classified gluttony as a deadly sin and depicted gluttons as suffering the torments of hell). But even those who think it is best that *no one* suffers may think it is better that those who are responsible suffer than those who are not. Given that our resources are limited, tough decisions have to be made and someone will always suffer. We may temper mercy with justice, since mercy must be tempered will we or no.

7.3 Obstacles to Responsibility

We temper mercy with justice by ensuring that the burdens of suffering fall on those who are guilty, given they must fall at all, only if it is indeed just that the guilty deserve to suffer to some extent; at least, that their right against such burdens is weaker than the same right of those who are not guilty. I defined basic desert in this kind of minimalist way in *Hard Luck*: to say that someone deserved negative consequences in virtue of their responsibility was not to say that it was good that they suffered, but only to say that such agents “no longer deserve the (full) protection of a right to which they would otherwise be entitled: a right against having their interests discounted in consequentialist calculations” (Levy 2011: 3). I went on to argue that in fact no one is responsible even in this minimalist sense: no one loses the right against having their interests discounted in consequentialist calculations (framed in the language of luck egalitarianism, I argued that no one was responsible for outcomes that are significantly due to luck because all our choices are shot through with luck). Here, I will set my own (idiosyncratic) views aside. Instead, I aim to show that even on more mainstream accounts of responsibility, there are large obstacles to confident attribution of responsibility in the case of lifestyle diseases.

On the consensus view of moral responsibility, for an agent to be morally responsible for an action or a state of affairs they must satisfy two independent conditions: a *control* condition and an *epistemic* condition.³ To say that the agent exercises (responsibility-level) control over an action or a state of affairs is to say

³ In *Hard Luck*, I argued that these two conditions were not independent (Mele (2010) has independently argued for a similar view). I set aside this idiosyncratic view here; in any case, it complicates the picture without changing it in its main features.

that it is appropriately sensitive to their decisions. While it is notoriously difficult to spell out what exactly is required for the possession of such control, the intuitive idea is clear enough. I control those things I can intentionally alter through my bodily movements (including those movements that consist in verbal behaviour). I am not responsible (in the relevant sense) for those things I cannot affect by acting. To say that the agent satisfies the epistemic condition is to say that she understands (at least implicitly) what the effects of her actions or omissions will be (or, in many versions, that she *ought* to understand). My putting salt in your coffee might be a (mildly) blameworthy act, if I know that the sugar bowl contains salt. But if I don't know that someone has switched the sugar for salt, I fail to satisfy the epistemic condition and don't deserve any blame.

Whether agents who come to need scarce health care resources as the result of their own actions satisfy these conditions has already received extensive discussion. There are serious worries that at least some classes of agent do not satisfy the control condition. For example, the fact that many of the most unhealthy behaviours involve addictive substances—with smoking and excessive consumption of alcohol accounting for the bulk of serious lifestyle-related morbidity and mortality—and that addiction impairs control is well-recognized. Even ordinary self-control problems may undermine responsibility attribution. Elsewhere, I have argued that the best explanation for the correlation between socio-economic status (SES) and unhealthy behaviours (and, correlatively, between SES and ill-health) is that lower SES tends to cause reduced self-control capacities and greater difficulty in resisting temptation (Levy 2019). Lower SES individuals are likely to face more stressors, which are known to reduce self-control, and live in neighbourhoods where temptations are harder to avoid, and so on. For those people most likely to suffer poor health induced by lifestyle, better choices with regard to non-addictive consumption and with regard to inactivity may be just as difficult as (or even more difficult than) better choices for addicts with regard to drugs. Indeed, it has been suggested that overeating may be harder to overcome than substance additions (Persson 2014). In addition, I have suggested that we all face problems with satisfying the epistemic conditions on responsibility for disease, because we live in epistemically polluted environments, in which discovering reliable information is difficult for those who lack the right kinds of networks or training (Levy 2018).

Even in the face of these complications, however, there's little doubt that agents who are *causally* responsible for their own ill-health regularly satisfy standard tests for moral responsibility too. The most influential test is that developed by John Martin Fischer, alone and together with Mark Ravizza (Fischer 2011; Fischer and Ravizza 2000). This tests probes agents' capacities: an agent is morally responsible for an outcome if she's capable of controlling whether it occurs (and she satisfies the epistemic condition with regard to it). She possesses the relevant capacities if she satisfies a counterfactual test: she *would* exercise her control over it given

reasons to do so. She must be *receptive* to these reasons (capable of recognizing some of them *as* reasons) and *reactive* to them (adjusting her behaviour in response to some of those she recognizes). Here's not the place to delve into the details of this test. Suffice it to note that even addicts satisfy them. Given a sufficient reason to refrain from smoking, even the heaviest smoker will do so, for a shorter or longer period of time. Despite the myths, the same is true of the heroin addict or the cocaine user. The proverbial police officer at the elbow will motivate every addict to refrain for the time being. But incentives need not be extraordinary to motivate abstinence. Cocaine addicts will refrain for an extended period of time in exchange for low-value vouchers for things like cinema tickets (Higgins et al. 1994; Lussier et al. 2006). In fact, most addicts seem to "mature out" of addiction on their own, and they seem to do in response to ordinary incentives (Heyman 2009). For example, landing a new job, entering a new relationship or having a baby all provide incentives that prove powerful for many addicts.

The counterfactual test for moral responsibility is designed with one-shot cases in mind, and for those purposes may be appropriate. In effect, the test asks how the agent would respond *were something significant at stake*, and when we're concerned with responsibility for crimes, something significant *is* at stake. We think it's perfectly reasonable to enquire after how agents would respond were they attentive and careful, because it's no excuse in contexts like this that we weren't paying attention. But in the kind of cases we're concerned with—smoking, say—much less is at stake on each token occasion, and a test that asks how the agent might respond if they were paying careful attention may not be the right one to apply. There's no doubt that agents who cause their own ill-health are sometimes careful and attentive, but it's an open question how often we satisfy conditions like this with regard to low stakes and often habitual behaviours. Special problems arise for the attribution of moral responsibility in contexts like this.

7.4 Responsibility for Patterns of Behaviour

Brown and Savulescu (2019) argue that holding agents responsible for their ill-health when it arises from lifestyle requires showing that the agents are *diachronically* responsible. Diachronic responsibility consists in being responsible for at least a minimum proportion of the repeated behaviours. As they point out, diachronic responsibility comes in degrees: an agent might satisfy the conditions on being responsible for *all* the instances of behaviour that contribute to ill-health, for the *majority*, or for some minimum threshold. The distinction between diachronic responsibility and what we might call one-shot responsibility is important. I will argue that we need some further distinctions beyond those that Brown and Savulescu introduce. Ill-health may arise from lifestyle in a variety of

ways, and different causal relations between lifestyle and ill-health have different implications for diachronic responsibility.

We can distinguish at least three kinds of cases where harm arises from repeated behaviour. The first set consists of cases that fit the *cumulative* model; smoking might be a good example. On the cumulative model, each instance of behaviour independently contributes to the risk of serious ill-health, though the contribution of each might be infinitesimal. When the behaviour is repeated many times (say a pack a day for twenty years), the cumulative risk is raised considerably simply through addition. The second model is the *agglomeration* model. On this model, no individual instance of the behaviour makes an independent contribution to the risk by itself, but excessive consumption significantly raises the risk. Perhaps some illegal drugs are like this: perhaps no single (carefully calibrated) dose of heroin raises the risk of ill-health for an otherwise healthy person, but many doses significantly raises the risk (sugar consumption might be a less contentious example). Substances that agglomerate to be harmful are those of which the old adage “the dose makes the poison” is true. On the agglomerative model, the risk is emergent from the behaviour. On the *stochastic* model, *most*—perhaps almost all—instances of the behaviour carry no negative consequences at all for the person, but each involves a risk of negative consequences. Unprotected sex with multiple partners might be like this: most partners will probably not carry a virus. But repeated behaviour raises the risk considerably.

These distinctions greatly complicate the attribution of responsibility: different patterns of causation may imply different degrees of moral responsibility for the same effects. In what follows, I will work through the different models. As we will see, there are important differences between them in what they imply for moral responsibility (and in particular how confident we have a right to be in attributing moral responsibility to agents for their own ill-health). I will *not* show that agents cannot be responsible for their own ill-health under any of these models. My ambitions are less grand: to sketch the complications that these models entail for the justified attribution of responsibility and to suggest that we often have much less right to confidence about such attributions than we might think.

7.4.1 The Cumulative Model

Let’s begin with the cumulative model. On this model, each instance of behaviour carries some independent risk, though it is typically low (if it were not low, most agents would not be willing to run the risk: the fact that smoking *this* cigarette is almost risk-free for me is a necessary condition of my smoking it). Are agents responsible for resulting ill-health?

Once a behaviour is established, it may become habitual. Smoking may easily become habitual because it is widely accepted (or at least was until recently) and

doesn't take a great deal of effort in preparation (in these ways, it is quite unlike, say, heroin use, which is unlikely to become habitual). Habitual behaviours are subject to a relatively low degree of control. The habit system may render behaviour "unintentional, robotic, perhaps even unconscious" (Redish et al. 2008: 424). Of course, behaviours that are habitual are not subject to a low degree of control all the time. Smokers give up smoking, obviously. Just as obviously, their behaviour is reasons-responsive (in the sense made famous by Fischer and Ravizza 2000): given sufficient reasons to inhibit the behaviour, the agent will succeed for a shorter or longer period of time. Nevertheless, very many of the behaviours that cumulatively give rise to the risk exhibit a low degree of control.

We don't blame agents for running small risks (the agent who takes a short-cut home knowing there's a one in a million chance of being hit by lightning surely takes a reasonable chance and isn't held responsible if she's subsequently struck). A simplistic framework for attributing responsibility to agents in cases like this would simply add up the token instances on which the agent exercised control for their behaviour and ask whether the total risk arising from those instances combined is sufficient to hold the agent responsible. Given the behaviour is habitual, it's likely that such instances are not sufficiently numerous to underwrite attribution of responsibility.

However, the simplistic framework is too simple. In attributing responsibility, we shouldn't treat these decisions to smoke (or to refrain from smoking for some period of time) independently of the agent's other behaviours. The agent may not have exercised responsibility-level control over many token behaviours that together raise the risk to some considerable level, but she's aware that she engages in them, and she ought to take them into account when she makes her attentive decisions. Analogously, I ought to take the behaviour of other agents (over whom I exercise no control) into account in making decisions about how to behave; if I vandalize a painting together with other people, for example, I'm not excused for causing damage on the grounds that my action wasn't sufficient by itself to cause it.

Nevertheless, the fact that control does fluctuate over time and most of the time I may not exercise responsibility-level control over my behaviour should surely reduce my responsibility for the outcome, relative to having caused it in a fully controlled fashion (analogously, the fact that my contribution to the damage wasn't sufficient to cause it might reduce my responsibility for the outcome). Moreover, there are other reasons to think that my responsibility might be reduced on the cumulative model (and on the agglomeration model too) in some instances.

Consider a case like this: an agent causes ill-health by excessive eating. Instances of eating come in two varieties for her. Some feature fine food that she savours; others involve habitual and inattentive gulping of junk for stress relief. The first kind is important to her and are among her most significant pleasures. The second

she barely notices. She exercises responsibility-level control only over the first kind, since she is sufficiently attentive only when she savours her food. While there are grounds for thinking she ought to refrain from her controlled eating, in asking her to refrain we ask her to forego a pleasure that is significant to her. Even if she values good health more than these pleasures, their importance to her and the place they occupy in her life seem to constitute grounds for mitigation.

Again, I don't take these considerations to show that she's not responsible. Rather, I take them to complicate attributions of responsibility and to give us reason to be cautious in making such attributions. Of course, we might hope to set all these complications aside by abstracting from the question whether or when the agent exercises control over her behaviour and asking instead about her acquisition of a habit. I set this question aside for the moment.

7.4.2 The Agglomeration Model

I turn now to the agglomeration model. On this model, a risk arises only when the number of behaviours passes a certain threshold. Prior to that threshold, there is no (or no measurable) risk at all. It seems to follow that on such models, an agent might be fully responsible for a great many of her behaviours, but not for a significantly raised risk of ill-health because she is responsible for some number of token behaviours insufficient to cross the threshold required for risk. If that's right, then the kinds of considerations introduced above—about how control fluctuates over time and how the epistemic condition for habitual behaviours may be harder to satisfy than at first seems—may entail that we should not be confident that people are responsible for agglomerative risks.

We've already noted both that this simple way of thinking about risks is too simple, because we shouldn't treat token controlled behaviours as independent of uncontrolled behaviours (at least if the agent is aware of them), and that there may be grounds for mitigating responsibility nevertheless that carry over from the cumulative model. Here I want to note a reason for caution that seems to rise especially pointedly on the agglomeration model. On these models, and depending where the threshold for risk is, attributions of moral responsibility confront a significant worry from moral luck.

Moral luck cases arise when an agent seems responsible for an outcome, but differs from other possible or actual agents who are not responsible for the outcome only in respects that are shot through with luck. The standard illustration is reckless driving: an agent who drives recklessly and injures a pedestrian may differ from another reckless driver who did not cause harm only in experiencing bad luck. Since luck is not a basis for desert, there are strong grounds for treating the agents alike: either blaming both or excusing both (Levy 2016a). Such cases can arise on the agglomeration model.

Consider two agents who engages in behaviour liable to give rise to such harms. One ceases the behaviour just prior to passing the threshold at which there is a significant risk of harm; one continues just beyond that threshold. Is there a difference in their responsibility if they both subsequently develop ill-health? In the case of the first agent, the ill-health is (by stipulation) not a result of their behaviour. Because they ceased the behaviour prior to passing the risk threshold, the ill-health arose by bad luck instead. The other agent may indeed have caused their ill-health (though of course we cannot be sure; we are rarely in a position to be confident that had the person not engaged in certain behaviours, they would not have suffered the problem). But they engaged in the behaviour only slightly more than their counterpart. Moreover, the first is unlikely to have known that they stopped just prior to the threshold (in fact, no one has any clear idea where that threshold is). They did not satisfy the control or epistemic conditions with regard to the fact *that they stopped in time*. So if they escape responsibility, it is in virtue of a difference which is slight and which is little to their credit (both, recall, did stop before developing ill-health). The difference between the two agents seems to be a matter of luck. By hypothesis, the agent who didn't cross the threshold for substantial risk didn't cause their ill-health and therefore isn't liable to bear any additional burdens; since she differs from the other agent only in her luck, there are grounds for excusing *both*.

7.4.3 The Stochastic Model

Before turning to the question whether responsibility can nevertheless serve as tie-breaker between patients, let's briefly consider the stochastic model. On that model, recall, most instances of behaviour do not cause any harm at all but some are likely to cause serious harm. *Epistemically*, token acts may be indistinguishable, but metaphysically they're unlike. We might call this the Russian roulette model. Pointing a gun at one's head and pulling the trigger when the chamber is empty carries no risk, whereas when there's a bullet in the chamber it is almost always fatal. Russian roulette is a gamble because the two cases look identical to players. Russian roulette is one example of a behaviour that fits the stochastic model. Another might be unprotected sex with strangers: on most occasions there may be no risk of a sexually transmitted disease but on some (epistemically indistinguishable) cases, the risk might be high. Sharing needles might be another example.

One way to attribute responsibility on such a model might be as follows: what matters is whether the agent is responsible for the *specific* behaviour that actually caused the harm. Suppose an agent contracts a serious illness from sharing a needle, and the illness was caused by one particular occasion out of eighty times she shared needles. Whether she's responsible for the illness might depend on

whether she was responsible for *that* behaviour. If that's right, the stochastic model might be particularly demanding: an agent might be responsible for *almost all* her actions and fail to be responsible for her ill-health. This intuition would be particularly strong with regard to behaviours that are morally innocuous, taken one by one.

However, just as we have reasons not to treat token behaviours independently of one another on the agglomeration and the cumulative models, so we have reason to take actions that didn't result in any harm (directly, at any rate) into account in assessing responsibility on this model. Even though these actions didn't cause ill-health, they may have made the action that *did* cause ill-health more probable. If some of these directly harmless actions are free, then they might underwrite attribution of responsibility to the agent.

Actions may make future (free or unfree) behaviours more likely in the future in one or more of several ways. They may establish habits. They may result in altered assessments of risk ("if sharing needles turned out okay a dozen times, it's probably not very dangerous"). They may alter the environment in which the agent finds herself, and therefore change the temptations or the opportunities (freely installing Tinder on your phone today might make unsafe sex more likely tomorrow, when your self-control is low). Of course, causal responsibility for later unfree behaviours isn't sufficient for moral responsibility for those later behaviours: the agent must satisfy the epistemic condition with regard to them as well. That is, she must know (or be culpable for failing to know) that her behaviour now makes future behaviour more likely. I set this question aside until the next section, which deals with such "tracing" conditions as they pertain to all three models. For the moment, I note only that such assessments may be difficult and at least sometimes agents won't satisfy the epistemic tracing condition sufficiently for their earlier actions to underwrite responsibility for later behaviours.

Just like the agglomeration model, problems of moral luck can arise on the stochastic model. Two agents may have (for example) each shared needles with others on dozens of occasions. Agent A may contract hepatitis on one such occasion, whereas agent B was luckier and did not contract any serious disease. We may hold the tracing condition fixed for both (recall that the problem of moral luck does not require that there is an *actual* luckier counterpart; merely that a counterpart that differed only in their luck could exist⁴). Since the agents differ only in their degree of luck, it is difficult to see how one might deserve a greater burden than the other.

⁴ It is not sufficient that a counterpart might have existed who did not suffer the harm. There are constraints on the attribution of luck, and probability features, in complex ways, among them. If an event was likely, then the agent is not lucky to have been subject to it. The agent who is hit by a car while walking across the freeway is not unlucky, and the fact that a possible counterpart was not subject to this harm does not provide her with an excuse. See Levy (2011) for discussion.

7.5 Tracing Responsibility

We sometimes want to hold agents responsible for later behaviours in virtue of their previous free actions (Vargas 2005). For example, we may want to blame an addict for her habit in virtue of their actions *prior to* the behaviour becoming habitual—and hence when the low degree of control characteristic of habits was not at issue. Such attributions are subject to a demanding epistemic condition. The epistemic condition entails, among other things, that agents must have some kind of grasp of what is at stake in their behaviour. They must understand that they are at significant risk of developing that pattern of behaviour *and* that if they develop that pattern of behaviour they will run a significant risk of subsequent ill-effects. These conditions are harder to satisfy than is often realized. People may overestimate their resistance to developing a habit or becoming addicted. Overheated drug rhetoric may contribute to agents possessing an inaccurate sense of the degree to which addictions are hard to shake: when they learn that their peers have tried cocaine or heroin without becoming addicted, they may think that popular rhetoric is false and may conclude that the risk is low.

Of course, habits often develop when the person is young, and may have difficulty in assessing risks. A smoking habit developed as a teenager will make later life abstinence much more difficult. Young people *also* have weaker control systems: they may fail to satisfy (or fully satisfy) both epistemic and control conditions. While a great many of the instances of behaviours that together give rise to a significant risk of ill-health are sufficiently controlled and sufficiently knowing to count as morally responsible, we often lack grounds for confidence that the agent is responsible for a sufficient number to count as a responsible for the ill-health should it eventuate.

There are also questions of moral luck that arise with regard to the epistemic tracing condition. Agents A and B may have fulfilled the tracing condition to the same extent, with only one developing a harmful habit (perhaps as a result of a genetic vulnerability they could not have known about). Most people who try hard drugs do not in fact become addicts, after all. Since the agents do not differ in the degree to which they satisfy the tracing condition, it seems unfair to hold them responsible to different degrees.

It is also worth emphasizing that some harmful habits develop out of behaviours that are difficult to avoid, or that are innocuous. Early instances of the behaviour might be prudentially neutral at worst; indeed, some instances of some behaviour may be prudentially *required*. The clearest example here is eating. Because eating is prudentially required, it has features that make it peculiarly difficult to control. Eating is unlike smoking or alcohol consumption in that it is not possible to treat the problem by abstinence. The rule “never smoke again” is maximally clear: the contexts to which it is relevant and the behaviours

that it prescribes (or proscribes) in these contexts are quite precisely specified. But rules like “do not overeat” are far from clear in their application and require a great deal of interpretation. How much is too much? How should episodes of consumption be balanced against one another and against other activities (e.g., going for a walk)? Is it permissible to eat the birthday cake *now* (perhaps required: my friend’s feelings will be hurt if I don’t accept) and skip breakfast tomorrow to make up for it? The intrinsic vagueness in dietary rules may entail difficulties in holding agents responsible for the development of harmful habits.

There are also difficulties arising with regard to satisfaction of the epistemic condition in this domain. Some dieticians argue that we should never see some foods as forbidden entirely, because that leaves the dieter feeling deprived and makes them more vulnerable to lapses (which may be seen as devastating failures). But the knowledge that no food is forbidden is itself open to abuse, since once again it is intrinsically difficult to calculate how much is okay. It is cognitively demanding to calculate how much we ought to eat, especially in contemporary environments where food is plentiful and omnipresent. Moreover, the fact that these rules require interpretation opens the way for self-deception in their application (Ainslie 2001; Levy 2016b). All these facts may entail that agents’ responsibility for overeating is lower than we might have thought. The agents in question failed at a task that is more difficult than the one facing those who do not experience constant temptation and a subsequent need to engage in cognitively demanding tasks (see Brown, this volume, for further discussion of the complexities in assessing the degree to which agents satisfy the tracing condition).

7.6 Responsibility as a Tie-Breaker

My aim, in the foregoing, has been both to shed light on the structure of behaviour that causes ill-health (in typical cases) and also thereby to suggest we have weaker grounds for concluding that agents are responsible for such ill-health than we might have thought. As we saw at the beginning of this paper, however, there are important differences between responsibility in the health context and responsibility in (say) the context of criminal law.

In the criminal law context, the standard of proof is reasonable doubt. If we have grounds for doubting either that the agent actually committed to crime or that she was responsible for committing it, we should acquit. But we may think we need to satisfy a lower standard of proof in the context of the allocation of scarce resources. Setting consequential considerations aside (as we may often appropriately do in the context of criminal law; in any case, such considerations are irrelevant to questions of basic desert by stipulation), no one is wronged by a

guilty agent going free.⁵ It is not necessary that a burden fall on anyone. In the resource allocation context, however, matters are crucially different. Resources expended one way are not available for expenditure in others. An agent who is more clearly not responsible for her ill-health may therefore have a claim on these resources that is stronger than another who is less clearly absolved of responsibility. It seems that we may therefore appropriately invoke responsibility, or a reasonable suspicion of responsibility, as a tie-breaker between agents. Other things being equal (each needs the resource to the same degree; the expected benefits to each are more or less identical; QALYs are equalized, and so on), we might prefer to allocate resources to agents who are more clearly absolved of responsibility.

The tie-breaker argument is surely the strongest available to the defender of moral responsibility in the allocation of scarce resources (though it is likely limited in its scope: there may be few cases in which we face a choice between agents who are roughly equal in the expected utility they might derive from a treatment). Those who invoke it may accept that they do not have the right to confidence that a particular agent is morally responsible for their ill-health but note—plausibly—that they have a better reason to think that the agent is morally responsible for their ill-health than another who needs the resource equally badly. If we accept such an argument, however, it will be at the cost of witting unfairness. That is, if we use responsibility as a tie-breaker, we do so in the knowledge that it will rule some people responsible who are not (and perhaps vice versa). That's a moral cost to us and to the individuals wrongly stigmatized as responsible, and it's an open question whether we should pay this cost.

Of course, we must accept that whatever we do, some people will be denied those resources to which they have an equal right. That's a sad entailment of the fact that need exceeds supply. We must allocate resources so some people are not treated as they deserve. But that we should run the additional risk of stigmatizing some in the process surely should not be accepted without further argument. We may prefer to allocate resources on other grounds: a lottery, or a queuing system, for example. While I won't argue for that claim here, the fact that the capacities and opportunities to escape blame are *themselves* unfairly distributed (Levy 2019) entails that when we use these kinds of devices we treat people *more* fairly than when we are sensitive to considerations of desert.

7.7 Conclusion

My aim in this paper has been to draw attention to unappreciated difficulties in assessing the extent to which people may be responsible for their own ill-health.

⁵ Or at least no one *need* be wronged. Some deontologists may think that if a guilty person goes free, their victims are wronged. But some crimes are victimless, and some victims may be dead. These kinds of effects are potentially important but not essential to the question of basic desert.

These difficulties are unappreciated because ethicists have paid little attention to diachronic responsibility and how it differs from one-shot cases. Further, they have not noticed how different kinds of causal relations between token behaviours and subsequent ill-health entail different sets of complications in the assessment of responsibility. I have not suggested that agents are not responsible for their ill-health under any of the models I have outlined. Rather, I have attempted to delineate the kinds of issues that confront us in attributing responsibility.

Firmer conclusions on when agents are responsible, under any of these models, must await further work which moves beyond the sketch of the landscape I have provided here and focuses on particular kinds of agents in particular contexts. In the meantime, I hope I have not only provided conceptual tools for such further work but also led to us seeing that attributions of responsibility of responsibility are often premature. We are not entitled to the confidence with which many of us blame agents for their health-related needs.

References

- Ainslie, G. (2001), *Breakdown of Will* (Cambridge University Press).
- Brown, RCH., Maslen, H., and Savulescu, J. (2019), “Against Moral Responsibilisation of Health: Prudential Responsibility and Health Promotion”, in *Public Health Ethics* 12: 114–29.
- Brown, RCH. and Savulescu, J. (2019), “Responsibility in Healthcare Across Time and Agents”, in *Journal of Medical Ethics* 45/10: 636–44.
- Caruso, GD. (2021), *Rejecting Retributivism: Free Will, Punishment, and Criminal Justice* (Cambridge University Press).
- Cohen, GA. (2011), *On the Currency of Egalitarian Justice, and Other Essays in Political Philosophy* (Princeton University Press).
- Davis, N. (2018), “Four in 10 cancer cases could be prevented by lifestyle changes”, in *Guardian*. 23 March 2018. <https://www.theguardian.com/society/2018/mar/23/four-in-10-cancer-cases-could-be-prevented-by-lifestyle-changes>
- Di Forti, M., Sallis, H., Allegri, F., et al. (2014), “Daily Use, Especially of High-Potency Cannabis, Drives the Earlier Onset of Psychosis in Cannabis Users”, in *Schizophrenia Bulletin* 40: 1509–17. <https://doi.org/10.1093/schbul/sbt181>
- Fischer, JM. (2011), *Deep Control: Essays on Free Will and Value* (Oxford University Press).
- Fischer, JM., and Ravizza, M. (2000), *Responsibility and Control: A Theory of Moral Responsibility* (Cambridge University Press).
- Heyman, GM. (2009), *Addiction: A Disorder of Choice* (Harvard University Press).
- Higgins, ST., Budney, AJ., Bickel, WK., et al. (1994), “Incentives Improve Outcome in Outpatient Behavioral Treatment of Cocaine Dependence”, in *Archives of General Psychiatry* 51: 568–76.

- Levy, N. (2011), *Hard Luck: How Luck Undermines Free Will and Moral Responsibility* (Oxford University Press).
- Levy, N. (2016a), “Dissolving the Puzzle of Resultant Moral Luck”, in *Review of Philosophy and Psychology* 7: 127–39.
- Levy, N. (2016b), “‘My Name is Joe and I’m an Alcoholic’: Addiction, Self-Knowledge and the Dangers of Rationalism”, in *Mind & Language* 31: 265–76.
- Levy, N. (2018), “Taking Responsibility for Health in an Epistemically Polluted Environment”, in *Theoretical Medicine and Bioethics* 39: 123–41.
- Levy, N. (2019), “Taking Responsibility for Responsibility,” in *Public Health Ethics* 12/2: 103–13.
- Lussier, JP., Heil, SH., Mongeon, JA., Badger, GJ., and Higgins, ST. (2006), “A Meta-Analysis of Voucher-Based Reinforcement Therapy for Substance Use Disorders,” in *Addiction* 101: 192–203.
- McKie, J., Singer, P., Kuhse, H., and Richardson, J. (1998), *The Allocation of Health Care Resources: An Ethical Evaluation of the “QALY” Approach* (Routledge).
- Mele, A. (2010), “Moral Responsibility for Actions: Epistemic and Freedom Conditions”, in *Philosophical Explorations* 13: 101–11.
- Neumann, PJ., Cohen, JT., and Weinstein, MC. (2014), “Updating Cost-Effectiveness: The Curious Resilience of the \$50,000-per-QALY Threshold”, in *New England Journal of Medicine* 371: 796–7.
- Pereboom, D. (2014), *Free Will, Agency, and Meaning in Life* (Oxford University Press).
- Persson, K. (2014), “Why Bariatric Surgery Should be Given High Priority: An Argument from Law and Morality”, in *Health Care Analysis* 22: 305–24.
- Redish, AD., Jensen, S., and Johnson, A. (2008), “A Unified Framework for Addiction: Vulnerabilities in the Decision Process”, in *Behavioral and Brain Sciences* 31: 415–37.
- Segall, S. (2010), *Health, Luck, and Justice* (Princeton University Press).
- Shaw, E., Pereboom, D., and Caruso, GD. (2019), *Free Will Skepticism in Law and Society: Challenging Retributive Justice* (Cambridge University Press).
- Vargas, M. (2005), “The Trouble with Tracing”, in *Midwest Studies in Philosophy* 29: 269–90.
- Volkow, ND., Swanson, JM., Evins, AE., et al. (2016), “Effects of Cannabis Use on Human Behavior, Including Cognition, Motivation, and Psychosis: A Review”, *JAMA Psychiatry* 73: 292–7.
- Yoon, PW., Bastian, B., Anderson, RN., et al. (2014), “Potentially Preventable Deaths from the Five Leading Causes of Death: United States, 2008–2010”, *Morbidity and Mortality Weekly Report* 63: 369–74.