

Could've known better

ALEXANDER GREENBERG ^{1,2}

¹University of Southampton, UK and ²Balliol College, Oxford, UK

Could you have taken precautions against a risk you were unaware of? This question lies at the heart of debates in ethics and legal philosophy concerning whether it's justifiable to blame or punish those who cause harm inadvertently or out of ignorance. But the question is crucially ambiguous, depending on what is understood to be inside or outside the scope of the 'could'. And this ambiguity undermines a number of arguments purporting to show that inadvertent wrongdoers cannot justifiably be blamed or punished. While not all opposition to blaming or punishing inadvertent wrongdoers rests on this ambiguity, some certainly does. And getting clear on this ambiguity is important if we're to sort good arguments against blaming and punishing inadvertent wrongdoers—if there are any—from bad ones.

Keywords: responsibility; blame; negligence; culpable ignorance; abilities.

I. Introduction

One of the most controversial questions about responsibility concerns whether people can be justifiably blamed or punished for things they do inadvertently. This question is asked by moral philosophers discussing when ignorance provides an excuse. If you poison your guest because you didn't realize there was arsenic in the sugar, in what circumstances does such ignorance provide an excuse and in what circumstances is such ignorance culpable (Rosen 2004: 299)? Similar questions are raised about unwitting omissions. Can you be justifiably blamed for failing to buy the milk your spouse asked you to pick up from the shops if you simply forgot to do so (Clarke 2014: 109)? Can you be justifiably blamed if you leave a dog, or a child, in a hot car to die of heat exhaustion because you forgot how long you had left them there (Sher 2009: 24)?

Correspondence to: Alexander Greenberg, a.greenberg@soton.ac.uk

The same question is asked by legal philosophers debating the justifiability of criminal negligence liability. Negligence, alongside recklessness and intention, is one of Anglo-American criminal law's main fault elements. Negligence consists in a criminal defendant's failure to take precautions against a risk of harm, precautions which a reasonable person in the defendant's situation would have taken. Negligence does not require any mental states on the part of the defendant, unlike other fault elements like intention and recklessness, the latter of which is standardly defined as involving awareness of risk. Someone can fail to take precautions against a risk of harm that a reasonable person would, and thus be negligent, even if unaware of the risk, and thus unaware of the need to take precautions against it. This is illustrated by the landmark English gross negligence manslaughter case of *Adomako*. Adomako was an anaesthetist who caused the death of a patient in his care during an operation because he failed to notice that an oxygen tube had come loose, leading to cardiac arrest. After an alarm sounded on the blood pressure monitor, Adomako spent several minutes trying to figure out what had happened, but did not check whether the tube was still connected. One expert witness described the standard of care as 'abysmal' while another 'stated that in his view a competent anaesthetist should have recognized the signs of disconnection within 15 seconds'.¹ Now Adomako clearly was aware that *something* was wrong. But he was unaware that his course of action—specifically, his failure to check whether the oxygen tube had come loose—created a risk of death. But Adomako's unawareness was no obstacle to him counting as *negligently* causing his patient's death. Is criminal liability—and the prospect of state punishment—justified for someone like Adomako who is unaware of the harm he is causing?

A central question in these debates is whether those who cause harm inadvertently have the requisite *control* over their conduct to count as responsible for the harm they cause. The claim that inadvertent wrongdoers lack control over their wrongdoing has figured in a number of sceptical arguments against the justifiability of blaming or punishing them.

The way of framing such concerns about control which I will focus on is one suggested by the following question: *Could an agent have taken precautions against a risk he was unaware of?* If an agent couldn't have taken precautions against such a risk, it seems manifestly unjust to blame or punish him for failing to do so. So it is crucial to figure out whether or not inadvertent agents are genuinely able to take precautions against the risks they inadvertently take.

However, the question of whether an agent could have taken precautions against a risk he is unaware of is ambiguous, depending on what is taken to be inside or outside the scope of 'could'. In this paper, I will suggest that this scope ambiguity undermines a number of sceptical arguments concerning

¹ *R v Adomako* [1995] 1 AC 171: 182.

culpable ignorance and criminal negligence. My claim is not that all concerns about whether inadvertent wrongdoers have the right kind of control over their conduct are dependent on exploiting this ambiguity. But I do aim to show that some arguments are guilty of this, and in so doing to put us in a better position to assess the claim that inadvertent wrongdoers fail to meet control conditions on responsibility.

II. A tempting sceptical argument

I will illustrate this scope ambiguity by outlining a tempting, but fallacious sceptical argument against blaming or punishing those who cause harms inadvertently. As we shall see, actual arguments guilty of this scope ambiguity are made in the moral and legal philosophy literatures. But it will be clearest to first state my own demonstrably fallacious argument.

Blaming, or punishing, those who cause harms inadvertently is blaming, or punishing, people for failing to take precautions against risks they were unaware of. This can make it seem inherently unfair to blame or punish those who cause harms inadvertently, if we appeal to the following three seemingly plausibly claims:

- (1) It's unfair to blame or punish someone for failing to take precautions against a risk if he couldn't have taken precautions against that risk.
- (2) Taking precautions against a risk requires being aware of that risk.
- (3) Someone couldn't have taken precautions against a risk if he was unaware of that risk.

Each claim on its face seems true. With regard to (1), if someone couldn't have taken precautions against a risk, it's not fair to blame or punish him if he fails to do so. (2) also seems true. I may have a growing lump in a lymph node which poses a risk to my health. But taking precautions against that risk—e.g. by seeing a doctor, getting a scan, seeking treatment, etc.—requires first being aware of the lump. For Adomako to take precautions against the risk that the disconnected oxygen tube poses to his patient's health, e.g. by reconnecting it, he must first become aware that the oxygen tube has become disconnected. And (3) seems to be entailed by (2). If taking precautions against a risk necessarily requires being aware of that risk, doesn't that mean that someone who is unaware of a risk couldn't have taken precautions against it?

And together these appear to give us reason to reject criminal negligence liability because (1) and (3) together seem to entail the following:

- (C) It's unfair to blame or punish someone for failing to take precautions against a risk if he was unaware of that risk.

Or so it seems. The argument is actually fallacious, because (3) features a crucial scope ambiguity introduced by ‘could’. To explain this, I’ll assume ‘could’ in the premises means *ability*, as that is plausibly the sense that matters for an argument about responsibility like this one (though I’ll say more about different meanings of ‘ability’ in the next section).²

On one reading of (3)—a *wide*-scope reading—(3) is definitely true:

(3^w) Someone couldn’t have (taken precautions against a risk if he was unaware of that risk).³

(3^w) is true and follows from (2). This reading says that we lack an ability to do a combination of two things together: (i) take precautions against a risk; (ii) be unaware of that risk. If (2) is true—i.e. if taking precautions against a risk requires being aware of it (a claim which I think is true)—this reading of (3) is undeniably true. One cannot do these two things together because taking precautions entails awareness. Taking precautions against a risk while at the same time being unaware of that risk would require both being aware and unaware of the same fact. And neither you nor anybody else—not even God—has the ability to do that. So given (2), this means there’s one sense—the wide-scope sense—in which you couldn’t have taken against a risk that you’re unaware of. That is, you can’t both be unaware of a risk and take precautions against that risk.

But consider the alternative, *narrow*-scope reading of (3):

(3ⁿ) (Someone couldn’t have taken precautions against a risk) if he was unaware of that risk.

(3ⁿ) says that someone who is unaware of a risk does not, while unaware, retain the ability to take precautions against this risk. Unlike (3^w), (3ⁿ) doesn’t follow from (2). In fact, (3ⁿ) is false. Taking precautions may require being aware of that risk, but that doesn’t entail that someone who is unaware of a risk lacks the ability to take precautions against it. Such a person may well retain the ability to take precautions if he still has the ability to *become aware* of the risk. There are many cases in which we could have become aware of risks which we weren’t in fact aware of. For example, you could have been aware of the risk of getting your foot caught in the gap between the train and the platform, something you didn’t realize there was a danger of, if you had paid more attention to the announcements over the Tannoy, and if paying more attention is something you could have done. You could have been aware

²To my knowledge, the only person who has pointed out that the scope of ‘could’ in such claims might matter in this context is Clarke (2023: 47). While Clarke does criticise some of my targets—notably Levy—he doesn’t argue their arguments trade off a scope ambiguity.

³There will be different ways of clarifying the logical structure of the two readings, but I trust that the way I have done this is clear.

that dinner's burning if you had figured out what that distinctive smell coming from the kitchen meant, and if figuring that out is something you could have done. And so on. And this means that, in such cases, we could have taken precautions against such risks. You could have known about the risks of getting your foot caught in the gap, and taken more care when alighting the train. You could have recognized that dinner's burning, and taken precautions against the risk that posed to your expensive new non-stick pan, if you had figured out what that smell coming from the kitchen meant. As things happened, you were too engrossed in watching television, with tragic consequences for your new non-stick pan.⁴

My suggestion is that in these cases and many others, because the inadvertent risk-takers retain the ability to become aware of a certain risk, they retain the ability to take precautions against that risk. In general, when an agent must φ in order to ψ , we often recognize an agent as having the ability to ψ even if he has not yet φ -ed because he has the ability to φ . For example, say I need to open the door to leave the house. Even though I haven't yet opened the door, it's still natural to regard me as having the ability to leave the house given that I can just open the door. My claim is that there is often a parallel relation between the ability to take precautions against a risk and the ability to advert to risk. The presence of the ability to advert to a risk often explains why an agent retains the ability to take precautions against a risk he is currently unaware of.

So while (3^w) is true, (3^n) is false, and only the former follows from (2). But this creates a problem for the tempting sceptical argument. This is because, although only (3^w) is true, the conclusion only follows if (3^n) is true.

Why is this the case? The fundamental logical point is that (3^w) doesn't entail that someone who is unaware of a risk thereby lacks the ability to take precautions against that risk. So even if we combine (3^w) with (1), it doesn't entail our conclusion.

A way of highlighting this, which may make the point more vivid, is to point out that if (3^w) did, together with (1), get us our conclusion, then obviously bad parallel arguments would also work. For example, it would have to be the case that the following argument is a good one:

- (I) It's unfair to blame or punish someone for failing to take precautions against a risk if she couldn't have taken precautions against that risk.

⁴My appeal to cognitive abilities here—e.g. abilities to figure things out, pay more attention, etc.—is influenced by Hart's answer to scepticism about criminal liability for negligence. Hart's central point is that we do not say that an agent 'couldn't have helped it' but taken a risk merely because she was unaware of the risk; we only say she 'couldn't have helped it' if she lacked the cognitive abilities to become aware of the risk because, e.g. she had a faulty memory, an inability to tell dangerous situations from safe ones, etc. (Hart 2008: 149–52). My framing of the argument in terms of a scope ambiguity can be thought of as a fuller explication of Hart's point.

- (II) Someone couldn't have (taken precautions against a risk if she deliberately didn't take precautions against that risk).
- (C) It's unfair to blame or punish someone for failing to take precautions against a risk if she deliberately didn't take precautions against that risk.

Premise (II) has the same structure as (3^w), and it's just as true as (3^w). But this argument is obviously absurd. The fact that someone deliberately didn't take precautions against a risk obviously doesn't mean that she lacked the ability to take precautions against that risk. So (I) and (II) don't entail the conclusion here. For the same reason, (1) and (3^w) don't entail the conclusion of our original argument. (1) and (3ⁿ) *would* entail that conclusion. But (3ⁿ) is false. So the tempting sceptical argument fails.

III. General and specific abilities

The point of setting out the tempting sceptical argument is to enable us to more easily identify the kind of scope ambiguity it exploits, in order to show how prominent arguments in the literature rest on such an ambiguity. But I first need to ward off a possible response concerning the meaning of 'ability'. So far I have left the meaning of 'ability' and 'can' intuitive. But 'ability' and 'can' can mean different things in different contexts. And those tempted by the tempting sceptical argument may appeal to different senses of 'ability' to try to respond to my objection.

In particular, we need to distinguish between *general* and *specific* abilities. Consider a French speaker with no other francophones to converse with. In one sense she is able to have a conversation in French, in another sense she is not. Consider a tennis player away from the court. In one sense he is able to play tennis, in another sense he is not. The solitary French speaker has a general but not a specific ability to have a conversation in French. The tennis player away from the court has a general but not a specific ability to play tennis.

This distinction is widely recognized, though sometimes different terms are used for it.⁵ The core of the distinction is that a specific ability to φ is an ability to φ in the specific circumstances the agent finds herself in. A general ability to φ is an ability to φ in a range of circumstances, but not necessarily in the circumstances the agent presently finds themselves in. An agent with a general ability to φ might lack a specific ability to φ in the circumstances she is in because in those circumstances she lacks the *opportunity* to φ , as our tennis player lacks the specific ability to play tennis because he's away from the

⁵See Honoré (1964); Kenny (1975: 131–44); Whittle (2010: 2–8); Way & Whiting (2016: 218–20).

court. An agent with a general ability to φ may also lack a specific ability to φ in certain circumstances because something *interferes* with the exercise of her ability in those circumstances. To illustrate with example of Ann Whittle's, 'Sally may have the ability to sing [i.e. the general ability] even though she freezes whenever her Aunt is present' (2010: 2).

Those tempted by the tempting sceptical argument might appeal to this distinction to respond to my objection—in particular, to respond to my claim that an agent who is unaware of a risk can retain the ability to take precautions against it, if she retains the ability to become aware of it. They might accept that someone unaware of a risk can have a general ability to take precautions against risks of that general kind, but deny that such a person has a specific ability to take precautions against *that* particular risk, in the circumstances she is in, and allege that the specific ability requires awareness.⁶ This would undermine my objection, because the specific ability to take precautions against a risk, rather than just the general ability, is plausibly required to justifiably hold someone responsible for failing to take such precautions.⁷ If I merely have the general ability to take precautions against a risk of a certain kind, but don't have the specific ability because I lack the opportunity or because something interferes, it doesn't seem fair to blame or punish me for failing to take precautions.

This response regards an agent's awareness of risk as constituting part of the circumstances required for them to exercise her ability to take precautions. The problem facing such a response is that whether or not an agent is aware of a risk doesn't plausibly look like part of the circumstances faced by the agent. We can spell this thought out in more detail by saying how a lack of awareness of a risk doesn't look like something which *interferes* with the exercise of one's ability to take precautions against that risk or something which removes one's *opportunity*.

Merely being unaware of a risk clearly isn't something which interferes with one's ability to take precautions against a risk. This is not to say that cases of inadvertence don't sometimes involve something interfering with the exercise of our abilities to take precautions against risk. To modify J.L. Austin's example, say I'm unaware that I'm stepping on your picnic spread because I'm distracted by a wasp (1956: 2). Something—the distracting presence of the wasp—interferes with me exercising my ability to take precautions against a risk—the risk of stepping on your picnic. But in a case like this, it's the wasp that's interfering with the exercise of this ability, by causing me to be unaware, rather than the unawareness as such.

⁶For this claim, see: (Nelkin and Rickless 2017: 127–28).

⁷As Kenny points out, the 'all-in can'—of ability plus opportunity—is 'the notion of "can" which is used in "can do otherwise" in discussions of responsibility' (1975: 141; cf. Whittle 2010: 10–2).

An opponent may suggest here that all cases of inadvertently caused harm in fact involve the same kind of interference as in the case of the wasp. After all, in all cases of inadvertent risk-taking, there will be *something* that explains why the agent was unaware. Won't that mean that there will always be something that counts as interfering with the exercise of the agent's ability to take precautions against the risk of harm?⁸

Such a general conclusion cannot follow from the mere fact that unawareness always has some explanation, because some explanations of unawareness are not aptly thought of as interference. In particular, unawareness of risk is often explained by some characteristic of the agent themselves, such as a psychological or character trait. Perhaps the agent has a hard time processing complex information quickly. Perhaps the agent is self-absorbed, so doesn't pay enough attention to where he is putting his feet. Such cases of inadvertent risk-taking may or may not be culpable. But they are not aptly thought of as cases where something 'interferes' with the agent's ability to take precautions against the risk. So we can make no general claim that cases of inadvertent risk-taking always involve interference.

What about opportunity? Does lack of awareness of a risk remove one's opportunity to take precautions against it? This also doesn't seem correct. Anthony Kenny unpacks the ability/opportunity distinction with the claim that '[a]n ability is something internal to an agent, and an opportunity is something external' (1975: 133). While Kenny admits that '[i]t is difficult to make this intuitive truth precise', if this criterion is more or less correct, then someone's awareness of a risk doesn't constitute part of the opportunity for exercising her ability to take precautions against it, because whether or not an agent is aware of a risk is clearly not something external to her.

Might one reject Kenny's claim that opportunities must be external to an agent, and allege that a lack of awareness of a risk can still mean one lacks the opportunity to take precautions against it? Kenny discusses a different case which will be helpful here, that of desire. Kenny claims that 'the mere lack of desire to do something, the mere presence of a desire to do the opposite, does not by itself remove the opportunity to do it' (1975: 133). This claim is plausible. If we recall our case of the tennis player, if he had his racket, an available court, and an opponent, but simply didn't want to play tennis today, and preferred to go to the beach, it would be incorrect to say that he thereby lacked the opportunity to play tennis.

Awareness looks similar to desire in this respect. Mere lack of awareness, by itself, doesn't look like it removes opportunity. Imagine that our tennis player is on holiday, and is asked by a friend if he can join her for a game of tennis on the hotel's court. His tennis racket is packed in his luggage, but he is unaware of this because it was packed by his wife. Does he lack the opportunity to play

⁸I am grateful to an anonymous reviewer for this response.

tennis? The claim I want to make is that *merely* lacking awareness of his racket doesn't suffice to mean he lacks the opportunity. If he had the (specific) ability to become aware that his racket was in his luggage—perhaps because he had the (specific) ability to pay attention to his wife telling him that she had packed it, rather than blithely ignoring her—then it seems wrong to say that he lacked the opportunity to play tennis with his friend.

So unawareness, *by itself*, doesn't look like it removes opportunity. What does look like it removes one's opportunity is *not being in a position to know* that one can exercise one's abilities. If our tennis player's wife had hidden a new tennis racket in his luggage as a surprise gift, and he had no reason to become aware or check his luggage while the offer to play with his friend still stood, then it does seem right to say he lacked the opportunity to play tennis. But what removes his opportunity to play is that he's not in a position to know he has a racket, not mere unawareness of it.

Similarly, if we return to the case of taking precautions against a certain risk, mere unawareness of the risk doesn't look like it removes one's opportunity to take precautions against it. Recall our earlier case in which you were unaware that dinner's burning and the risk that posed to your non-stick pan. If you had the (specific) ability to become aware of the risk—e.g. by figuring out what that distinctive smell from the kitchen meant—then it is incorrect to say you lacked the opportunity to take precautions against the risk to your non-stick pan.

We can thus conclude that whether or not an agent is aware of a risk is not plausibly considered part of the circumstances facing the agent. So the general/specific ability distinction does not provide an answer to my objection. The cases I've relied on to answer the tempting sceptical argument plausibly show that it's possible for an agent who is unaware of a risk to retain the *specific* ability to take precautions against that risk, if she has the specific ability to become aware of it. And whenever I speak of 'abilities' in what follows, I'll have specific abilities in mind.

IV. The guilty

We can now return to the point of setting out the tempting sceptical argument, which is to show how prominent arguments in legal and moral philosophy are guilty of the same kind of scope ambiguity it exploits. I'll demonstrate this with an example from each literature.

IV.1 Alexander and Ferzan on negligence

Our first example of an argument guilty of the scope ambiguity I've identified can be found in the following passage from Larry Alexander's argument

against the criminalization of negligence:

The negligent agent is ... not advertent to the antisocial quality of the contemplated act or omission. Although he can control whether he chooses so to act or refrain from acting, he is not advertent to the aspects of the act or omission that give him reason to choose otherwise. Although there are external reasons to choose otherwise, those reasons, not being occurrently internal, cannot motivate him. ... So although we can think about things we are not currently thinking about – if a reason to do so occurs to us – when a reason to advert to risks to which we are not presently advertent does not occur to us, we have no control over whether we advert to those risks. (Alexander 2014: 203)

Alexander is objecting to criminal liability for negligence on the grounds that the negligent agent's reasons not to act—'the antisocial quality of the contemplated act or omission'—cannot motivate him to act otherwise. In making this argument, Alexander is appealing to the claim, which is in itself correct, that being motivated by reasons requires awareness of those reasons.

However, the conclusion he makes from that claim—that the negligent agent cannot be motivated by the reasons there are to act otherwise—has the same kind of scope ambiguity we've identified in the tempting sceptical argument. There is one reading of Alexander's conclusion—a wide-scope reading—on which it is true:

(A^w) One cannot (be motivated by reasons to act otherwise if one is unaware of those reasons).

This is true because being motivated by a reason plausibly does require awareness of it. But there is another, narrow-scope reading of Alexander's conclusion:

(Aⁿ) (One cannot be motivated by reasons to act otherwise) if one is unaware of those reasons.

This reading of Alexander's conclusion is false. Despite the negligent agent's unawareness of the fact that constitutes his reason to act otherwise—the fact that his conduct carries a risk of harm—it's plausible that he nevertheless does retain the ability to be motivated by that fact *if* he retains the ability to become aware of it. And it's plausible that in many cases negligent agents do retain the ability to become aware of the risks attendant on their conduct if they're able to, e.g. pay more attention, think more about what they are doing, etc.

The problem for Alexander is that he needs (Aⁿ) rather than (A^w) to get any sceptical argument against the criminalization of negligence liability off the ground. We can see this if we think about the nature of the demand which criminalizing negligence places on those who cause harm inadvertently. Alexander's argument is clearly intended to show this demand is illegitimate.

But the demand implicit in criminalizing negligence is quite clearly *not* a demand to take precautions against the risks attendant on one's conduct *while remaining unaware of those risks*. That's obviously an illegitimate demand because it's one that's impossible to comply with, and the true reading of Alexander's conclusion—(A^w)—shows it to be illegitimate. But the real demand implicit in criminalizing negligence is a demand to take precautions against the risks attendant on one's conduct, risks which one may *happen to be unaware of*, but risks which one obviously would *become* aware of if one were to heed the demand and take the relevant precautions. And to show that to be an illegitimate demand by showing it is a demand we cannot be motivated to act on, Alexander needs (Aⁿ) to be true. But (Aⁿ) is false.

Now Alexander may respond here by appealing to his claim that 'we have no control over whether we advert' to risks we are not presently advertng to (and don't think we have reason to advert to). He may argue that because we lack control over whether we advert to risks, I cannot rely, in rejecting his argument, on the claim that someone who is unaware of a risk can retain the ability to become aware of that risk.

Is it true that we have no control over what we are aware of? Alexander does in fact argue for this claim, in joint work with Kimberly Ferzan, when they answer H.L.A. Hart's claim that criminal negligence liability can be justified if the defendant could have adverted to the risk of harm (Hart 2008: 149–57). Alexander and Ferzan dismiss Hart's claim as insufficient to ground criminal culpability for negligence:

If we take the actor at the time of the 'negligent' choice, with what he is conscious of and advertng to, his background beliefs, and so forth, then it is simply false that the actor 'could have' chosen differently in any sense that has normative bite. For although it may be true that the actor "could have" chosen differently in a sense relevant to the freewill-determinism issue, it is false that in that situation the actor had any internal reason to choose differently from the way he chose.

To have such a reason, an actor will have to advert to that to which he is not advertng. But one has no control at such moments over what one is advertng to or is conscious of: try thinking of what you are not thinking of, but should be! (Alexander and Ferzan 2009: 83)

Alexander and Ferzan are suggesting here that a negligent defendant would only have the kind of ability to choose differently which genuinely grounds culpability—the kind which has 'normative bite'—if he were able to 'advert to that which he is not advertng'. Their claim is that no one has such an ability, a claim they support by asking the reader to 'try thinking of what you are not thinking of, but should be!', the suggestion being that this is impossible thing to do.

However, as an argument for the claim that no one is able to advert to something he is not aware of, this is guilty of the exact same scope ambiguity

as the tempting sceptical argument. There is, again, one reading on which the claim is true, namely:

(AF^w) No one is able to (advert to something he is not aware of).

This is clearly true because advertent to something implies being aware of it, so no one has the ability to advert to something and at the same time be unaware of it. But again the alternative reading is what is required for the sceptical argument to go through, namely:

(AFⁿ) (No one is able to advert to something) he is not aware of.

Again, this reading is false. In many cases, we retain the ability to advert to—that is, become aware of—things we happen to be unaware of, as illustrated in previous examples. You retained the ability to advert to the risk of getting your foot stuck between the train and the platform because you could've paid more attention to the announcements. You retained the ability to advert to the risk of your new non-stick pan being ruined because you could've figured out what that smell from the kitchen meant. So (AFⁿ) is false. But, again, (AF^w), rather than (AFⁿ), is the claim Alexander and Ferzan need for their sceptical argument to work.

There are at least two responses Alexander and Ferzan may make at this point. First, they are likely to object that I have not shown that a negligent risk-taker has the ability to take precautions against risk *at the time at which he takes the risk*. They will point out that in the examples I've just given, the agents could have adverted to the relevant risks *in the past*. Once these agents have failed to pay enough attention or figure things out, they *now*, in the present, lack the ability to advert to these risks, and thus now lack the ability to take precautions against them. After all, what Alexander and Ferzan denied was that the negligent agent could have chosen differently 'if we take the actor *at the time of the "negligent" choice*' (my italics). Thus, in my previous example, when you negligently stepped off the train, Alexander and Ferzan will deny that *at that time* you retain the ability to take precautions against the risk of getting your foot stuck. Once you've failed to pay attention, and are thus unaware of the risk, it's too late.

This response appeals to a 'pinpointing strategy', one which focuses in on a particular time and claims that *that* is the relevant time at which we should assess the agent's ability to do otherwise. Such a response is natural because the pinpointing strategy is often legitimate. Say I've promised to pick up my mother from the hospital, but when I come to leave, a water main unexpectedly bursts outside my house resulting in my road being closed. The fact that I had the ability to pick my mother up *at an earlier time before the water main burst* isn't relevant to assessing my responsibility; it seems appropriate to assess what I could have done at the later time.

But the pinpointing strategy isn't always legitimate. This can be seen by looking at cases in which people deliberately make themselves unable to fulfil their obligations. Let's adapt an example from Walter Sinnott-Armstrong (1984: 252; cf. White 1975: 149). Say my friend has agreed to attend a concert with me, but is nowhere to be seen with the concert minutes away. When I worriedly call him, he tells me he can't make it in time. When asked whether something came up, he responds 'No. Nothing like that. I just decided not to leave in time. But given that happened, it's *now* the case that I can't make it. So you can't blame me'. This clearly is an unconvincing use of the pinpointing strategy, and for obvious reasons. My friend clearly *was* able *at a relevant time beforehand* to make it to the concert. And it's only the case that he is *now* unable to make it to the concert because he failed to do something—i.e. leave in time to make it—which he could and should have done.⁹

Using the pinpointing strategy to respond to my objection, I suggest, is illegitimate for the same reason as my friend's use of it is. You clearly did, at a relevant time beforehand, have the ability to take precautions against the risk of getting your foot stuck between the train and the platform because you had the ability to become aware of that risk. It may be true that you lack the ability to take precautions *at the time at which you step off the train*. But that is because you failed to advert to the risk—because you failed to pay attention to the announcements, failed to look before stepping off, etc.—which you could and should have done.¹⁰

So a response appealing to the pinpointing strategy fails. If it were to succeed in responding to my objection, it would overgeneralize to cases of deliberate wrongdoing that are undeniably culpable, such as the case of my friend who missed the concert. In both this kind of case and the cases of inadvertent risk-taking I've been discussing, what actually matters is that the agent had the ability to avoid wrongdoing across a relevant range of time before the act or omission we're blaming them for. And in both kinds of case, the agents I've been discussing have that ability.¹¹

We can now move on to the second response Alexander and Ferzan may make. They may stress that their point is not that we lack *any* control over what we are aware of, but that we lack the right *kind* of control. Perhaps their idea is that we lack the ability to become aware of risks we are unaware of *deliberately* or *by choice*. That may be why Alexander and Ferzan ask us to 'try' to think of what we are not thinking of.

⁹Such cases are often claimed by those who raise them to be counterexamples to 'ought' implies 'can'. My use of the example doesn't require that claim.

¹⁰It may be pointed out that my friend differs from this inadvertent agent because the former *chooses* to make himself unable to fulfil his obligations. This is true, but not a relevant difference unless Alexander and Ferzan are going to beg the question and assume only chosen wrongdoing is culpable from the get-go.

¹¹For a related objection to how Alexander and Ferzan pinpoint in on particular times, see Baron (2020: 78–79).

A move some philosophers will be attracted to here is to appeal to ‘tracing’, and suggest that we can in fact become aware of risks deliberately or by choice in some cases. Some of our cognitive abilities can be exercised deliberately. We can choose to pay attention to something, for instance. And we can also influence the exercise of our cognitive abilities by, for example, deliberately choosing to gather more evidence about risks, or to create some kind of reminder to check for risks.

I am sceptical that this is the right move to make at this point. I think the fact we can deliberately choose to perform such actions plays very little role, if any, in explaining responsibility for inadvertent wrongdoing. We may be able to exercise some of our cognitive abilities by choice or otherwise to deliberately influence their exercise. But this doesn’t do much to explain culpability for inadvertent wrongdoing. This is because failures to exercise our cognitive abilities—e.g. failures to pay enough attention—and failures to carry out such evidence-gathering actions—e.g. failures to ask if a patient is allergic—are typically inadvertent omissions, and so merely push the problem back.¹²

The right response here is instead to dig our heels in and question the presupposition that an ability to advert to risks *deliberately* or *by choice* is necessary for the negligent agent to be culpable. Why does it not suffice that the negligent agent is *able* to advert to the risks—even if he might not be able to do so deliberately, or if his failure to advert to risks was not deliberate?

One possible answer here would be to claim that if a negligent agent is unable to advert to a risk deliberately, he is then unable to take precautions against that risk deliberately. However, this claim just seems to be false. Even if we grant both (i) that culpability for failing to take precautions requires the ability to take precautions deliberately, and (ii) that the negligent agent is not able to advert to the risks deliberately, we can still maintain that the negligent agent remains able to take precautions against those risks deliberately. He does so because he retains the ability to become aware of the relevant risk. One would only have reason to deny that negligent agents retain the ability to take precautions deliberately if one required that the other abilities grounding that ability must themselves be abilities one can exercise deliberately. But this requirement is one we have no reason to accept, and some reason to reject. First, the requirement creates an obvious threat of regress. Secondly, there are clear counterexamples. Consider, for example, all the different abilities to move your body in certain ways that ground your ability to catch a ball. Can

¹²In this scepticism about the explanatory potential of tracing, I am in agreement with Alexander & Ferzan (2009: 79–81); see my (Greenberg 2021). I should note that my scepticism does not concern tracing ignorance back to *something* in the past. No one should object to that, and I just appealed to it in my answer to the first response. My scepticism concerns the explanatory significance of our ability to *deliberately* perform such ignorance-causing (or -avoiding) actions.

you exercise all of those underlying abilities deliberately? Clearly not. But you can nevertheless catch a ball deliberately.

So the question remains: Why doesn't it suffice for the negligent agent to be culpable that he is able, even if not deliberately, to advert to the risks? And it's not clear that a non-question begging answer—i.e. an answer that doesn't just assume that only deliberate or chosen wrongdoing is culpable—can be provided to this question. And until a non-question-begging answer can be provided to this question, it doesn't seem we have a plausible sceptical argument of the kind Alexander and Ferzan are aiming to provide, i.e. one that demonstrates the negligent agent is not culpable because he lacks the requisite kind of control over his conduct.

IV.2 Levy on culpable ignorance and luck

Our next example of a sceptical argument guilty of this scope ambiguity is given by Neil Levy, in his discussion of the extent to which ignorance excuses. A central kind of case discussed in the moral responsibility literature is the case of *moral* ignorance, such as the following commonly discussed case:

[C]onsider Mr. Potter, a powerful businessman who holds false moral views. He takes certain business practices—such as liquidating Bailey's Building and Loan and sticking it to the poor families of Bedford Falls—to be 'permissibly aggressive,' when in fact they're 'reprehensibly ruthless.' This leads him to do bad things, though he doesn't understand that he's acting badly, which means that he's acting out of a certain kind of ignorance. He's fully aware of all the circumstances, but he applies flawed normative principles or weightings and comes up with bad decisions. (FitzPatrick 2008: 599–600)

To say whether someone like Mr. Potter is culpable for what he does out of his moral ignorance, Levy claims that we need to answer the following question, suggested by William FitzPatrick:

R: What, if anything, could the agent reasonably (and hence fairly) have been expected to have done in the past to avoid or to remedy that ignorance? (FitzPatrick 2008: 603)

But while Levy agrees that this is the question to ask, he disagrees with how FitzPatrick answers it. FitzPatrick claims that it's often reasonable to expect agents like Mr. Potter to have done something to avoid or remedy their ignorance, if they have the capabilities and opportunities to do so and if their ignorance results from some vice, such as 'overconfidence, arrogance, dismissiveness, laziness, dogmaticism, incuriosity, self-indulgence, contempt, and so on' (2008: 605).

Levy claims it is not reasonable to expect agents like Mr. Potter to have done something to avoid or remedy their ignorance. 'So acting', Levy writes, 'is (at least typically, and in the relevant cases) not something that agents can do rationally and it is unreasonable for us to ask people to behave in ways that are not rational for them' (2011: 126–7).

Why is avoiding or remedying his ignorance not something someone like Mr. Potter can do rationally? It is here where Levy's argument trades off the scope ambiguity. Levy claims that what someone can do rationally—in the sense which is relevant the agent's responsibility—should be understood 'internalistically', by which he means that it's 'a function of what she takes her reasons to be' (2011: 127). So if we ask FitzPatrick's question R, and ask if it's reasonable to expect someone to avoid or remedy her ignorance, we can only answer affirmatively if the agent took herself to have reasons to do so. Levy explains why in the following passage:

Why should we understand R in an internalist manner? The answer is simple: it is only reasonable to demand that someone perform an action if performing that action is something they can do rationally; that is, by means of a reasoning procedure that operates over their beliefs and desires. But what agents can do rationally in this sense is a function of their internalist reasons. Suppose that there is a divergence between internalist and externalist reasons: what I have most reason to do, externalistically, is not something that I take myself to have any reason at all to do. In that case, if I do what I have most (externalist) reason to do, I do so not as a result of weighing reasons or any other reasoning procedure. I do so by chance, or through a glitch in my agency, or what have you. (Levy 2011: 128)

The crucial claim Levy is appealing to here—that someone cannot do something rationally if he is unaware of the reasons in support of it—has an initial ring of plausibility, because doing something rationally plausibly does involve awareness of the reasons in support of it. However, Levy's claim is ambiguous in the exact way I have identified. It is true on one reading, but false on another. And Levy needs the false reading for his sceptical argument to go through.

There is, as before, a wide-scope reading of Levy's claim:

(L^w) One cannot (perform an action rationally if one does not take there to be reasons in favour of that action).

This reading states that one cannot do two things together: (a) perform an action rationally; (b) not take there to be reasons in favour of that action. And this reading is plausibly true—and is what gives Levy's overall argument the ring of truth—because doing something rationally does look like it involves awareness of the reasons supporting the action. Say a son wants to poison his sick mother to hasten her demise and gain his inheritance more quickly. The son gives his mother what he takes to be poison, but what is in fact medicine that cures her. The son does what there is reason to do—cure his mother—but does not do so rationally, plausibly because that would require his taking there to be reason to cure his mother, as well as appreciating that he is curing her. And Levy's remarks about chance look apposite in relation to this reading

of his claim. It *is* just chance, or good luck, that the son did what there was reason to do, namely cure his mother.

But again there is another, narrow-scope reading of Levy's claim:

(Lⁿ) (One cannot perform an action rationally) if one does not take there to be reasons in favour of that action.

This reading says that if one doesn't take there to be reasons in favour of an action, then one does not retain the ability to perform that action rationally. And, again, this reading is false. One plausibly does retain the ability to perform the action rationally if one retains the ability to become aware of the reasons in favour of it. And in many cases we *do* retain the ability to become aware of reasons we aren't in fact aware of, if we are able to pay more attention, give more thought before acting, and so on.¹³ In my earlier example in which you let the food burn because you didn't realize what that smell coming from the kitchen indicated, you are unaware of a fact—the fact that dinner's burning—which constitutes a reason to take the pan off the heat. Yet because you retain the ability to become aware of that fact, you plausibly retain the ability to take the pan off the heat rationally. Likewise, despite Mr. Potter's unawareness of reasons to remedy his moral ignorance, if we flesh out the case in certain ways, he does look like he retains the ability to remedy his ignorance rationally. He plausibly retains this ability if he retains the ability to reflect on his own moral beliefs, to recognize that very many seemingly rational people disagree with him, to consider whether his moral beliefs might be influenced by self-interested bias, and so on.

If this is correct, it means that Levy's remarks about chance fall flat in relation to this narrow-scope reading of his claim. In the burning pan case, it's simply false to say that you could only perform the action the reasons in fact recommend—going into the kitchen and taking the pan off the heat—out of chance or a 'glitch in your agency'. It's equally false to say that Mr. Potter can only do what the reasons in fact recommend—i.e. change his business practices—out of chance or a 'glitch in his agency'.

Furthermore, as previously, it is the false narrow-scope reading that is required to get Levy's sceptical argument off the ground. Again we can see this if we think about the nature of an expectation to remedy one's ignorance, an expectation FitzPatrick thought fairly applied to morally ignorant agents and Levy did not. This expectation is not plausibly understood as a requirement to remedy one's ignorance while remaining unaware of the need to do so. The true wide-scope reading of Levy's claim does show that requirement to be unfair, but it's clearly not how we should understand the relevant expectation. The expectation is a requirement to remedy one's ignorance, which will,

¹³This point is made against Levy by Clarke (2017: 249–50), though he doesn't highlight Levy's argument is guilty of a scope ambiguity.

if heeded, involve becoming aware of the need to do so. The narrow-scope reading of Levy's claim is required to show that requirement is unfair. But the narrow-scope reading is false. So Levy's sceptical argument fails.

As an aside, it's worth pointing out that Douglas Husak, in his argument that ignorance of law should usually provide an excuse, relies on this argument of Levy's at a crucial point (Husak 2016: 153–5). So if Levy's argument fails, Husak's does too. This especially clear in Alex Sarch's schematization of Husak's argument, which includes the following premise: 'You cannot be motivated by, reason from, or be guided by a fact that you don't believe or whose truth you don't have any credence in whatsoever' (Sarch 2022: 321). This premise is ambiguous in the exact same way as Levy's claim was, and so Husak's argument fails for the same reason.

As before, there is response Levy is likely to make here, one that parallels a response already discussed. Levy will be apt to focus on the abilities I have claimed that ignorant agents retain, such your ability to figure out what that smell coming from the kitchen indicates, and Mr. Potter's ability to reflect on whether the fact that many disagree with him gives him reason to doubt his moral views. These are *cognitive* abilities, abilities to form certain beliefs, make certain inferences, assess one's evidence, recognize what reasons there are, etc. Levy will claim that we don't have the right kind of control over the exercise of these abilities. Levy writes that 'agents are never responsible for their beliefs, because they do not exercise freedom-level control over them' (2011: 112). He likewise claims that we don't have control over what reasons we recognize, suggesting that which considerations strike us as reasons at the time of action is a matter of luck or chance (2011: 90–1).

There is something wrong about what Levy says here. Specifically, it seems false to say that which considerations strike us as reasons is a matter of luck or chance, at least insofar as the ordinary concepts of luck or chance are concerned. Which considerations strike us as reasons depends, at least in large part, on the cognitive abilities we possess, and our epistemic virtues (or vices). These are factors that aren't appropriately described as being a matter of luck or chance.

Nevertheless, there is something importantly right about this response. Our possession of cognitive abilities may mean we have a *kind* of control over what we think and are aware of—and this kind of control may help distinguish someone like Mr. Potter from people with severe learning disabilities and people suffering from psychosis. But we still lack the kind of control over the exercise of our cognitive abilities that we have over our actions. Mr. Potter can raise his hand, or sign a cheque, *at will* or *by choice*, whereas he cannot, it seems, come to appreciate the error of his ways at will or by choice. We don't have the same *kind* of control over what we think as over what we do.

Is my objection to Levy undermined by the difference between the kind of control we have over what we think and what we do? It is not, for the same

kind of reasons as those I gave when answering the parallel response from Alexander and Ferzan. If Levy insists that culpability for wrongdoing requires the ability to avoid that wrongdoing rationally *and* by choice, I am happy to grant that claim. My claim is just that Mr. Potter retains that ability—he retains the ability to, say, change his business practices rationally and by choice—if he retains the cognitive ability to become aware of reasons to do so. The fact that this cognitive ability is not an ability he can exercise by choice is not an obstacle to this claim. It would only be an obstacle to my claim if Levy were to insist that each of the other abilities that ground or explain Mr. Potter's ability to avoid his wrongdoing rationally and by choice must also be abilities that he can exercise by choice. But, as before, we should reject this claim. It's simply not the case that all abilities that explain why I can choose to do something must be abilities I can exercise by choice.¹⁴ Recall my previous example of the abilities you have to move your hand in different ways that ground your ability to catch a ball. You cannot exercise all of those abilities at will, but you can catch a ball at will. To add another example, if you're not red-green colourblind, you possess the ability to bring me a red flower by choice or at will. But this ability is grounded, in part, by an ability to perceptually discriminate the colour red. That is not an ability you can exercise by choice or at will. So we can grant Levy the claim that the ability to avoid wrongdoing rationally and by choice is necessary for culpability. This claim would not show that someone morally ignorant, like Mr. Potter, lacks the ability to avoid his wrongdoing rationally and by choice. There is no good argument of the kind that Levy makes that agents like Mr. Potter lack that ability.

Now it may be objected here that the examples just appealed to are importantly different from the cognitive abilities which I am claiming inadvertent agents have. It may be objected that the ability to perceptually discriminate the colour red is an ability one cannot but exercise (if one possesses it and the relevant stimulus is present). That is not the case with Mr. Potter's cognitive ability to come to recognize the error of his ways, which he might or might not exercise.¹⁵ While this difference is a genuine one, it does not matter to the point I am making with the perceptual ability example. I am not claiming that the cognitive abilities that Mr. Potter possesses—or the cognitive abilities that negligent risk-takers possess—are exactly analogous to the ability to perceptually discriminate the colour red. They are not. The point of the example is just to show why we should not require that all the abilities that ground an ability to do something by choice must be abilities that can also be exercised by choice. And the example does show that to be the case, even given

¹⁴As Achs points out, 'it isn't true that, in general, having voluntary control over something requires having voluntary control over whether one meets the preconditions on which that control depends' (2020: 3707, fn. 24).

¹⁵I am grateful to an anonymous reviewer for this point.

the genuine differences it has from the other cognitive abilities I have been discussing.

V. Conclusion

I have shown how two prominent sceptical arguments concerning the justifiability of blaming and punishing inadvertent agents given in the legal and moral philosophy literatures rest on the scope ambiguity identified in the tempting sceptical argument. Insofar as scepticism about whether inadvertent agents have the control required for responsibility rests on such arguments, such scepticism is unfounded. Now I have not claimed that all doubts about whether inadvertent agents meet control conditions on responsibility rest on such a scope ambiguity. But I have shown that some do. If sceptics about blaming or punishing inadvertent wrongdoers still want to maintain that such agents lack the requisite control over their conduct, their task is to formulate an argument which shows that, but which does not rest on this scope ambiguity. I am sure it is possible to formulate a such a sceptical argument. Whether any such argument is plausible is another matter.¹⁶

References

- Achs, R. (2020) 'Blameworthiness and Constitutive Control', *Philosophical Studies*, 177: 3695–715.
- Alexander, L. (2014) 'Hart and Punishment for Negligence', in C. Pulman (ed.) *Hart on Responsibility*, pp. 195–205. London: Palgrave Macmillan.
- Alexander, L. and Ferzan, K. K. (2009) *Crime and Culpability: A Theory of Criminal Law*. Cambridge: CUP.
- Austin, J. L. (1956) 'A Plea for Excuses', *Proceedings of the Aristotelian Society*, 57: 1–30.
- Baron, M. (2020) 'Negligence, Mens Rea, and What We Want the Element of Mens Rea to Provide', *Criminal Law and Philosophy*, 14: 69–89.
- Clarke, R. (2014) *Omissions: Agency, Metaphysics, and Responsibility*. Oxford: OUP.
- Clarke, R. (2017) 'Ignorance, Revision, and Commonsense', in P. Robichaud and J.W. Wieland (eds) *Responsibility: The Epistemic Condition*, pp. 233–51. Oxford: OUP.
- Clarke, R. (2023) 'I Didn't Think of That', *Philosophical Issues*, 33: 45–57.
- FitzPatrick, W. J. (2008) 'Moral Responsibility and Normative Ignorance: Answering a New Skeptical Challenge', *Ethics*, 118: 589–613.
- Greenberg, A. (2021) 'Why Criminal Responsibility for Negligence Cannot Be Indirect', *The Cambridge Law Journal*, 80: 489–514.
- Hart, H. L. A. (2008) *Punishment and Responsibility: Essays in the Philosophy of Law*, 2nd edn. Oxford: OUP.

¹⁶For helpful comments on previous drafts, I would like to thank Rachel Achs, Maria Alvarez, Lucy Campbell, John Hyman, Conor McHugh, Findlay Stark, and two anonymous referees for *The Philosophical Quarterly*, as well as audiences in Cambridge, Freiburg, and Southampton. Special thanks to Lucy Campbell and Alex Sarch for discussions that prompted the idea for this paper, as well as to Mark Dsouza and Liat Levanon for letting me try out an early version of these arguments in the *Criminal Justice Theory Blog*.

- Honoré, A. M. (1964) 'Can and Can't', *Mind*, 73: 463–79.
- Husak, D. (2016) *Ignorance of Law: A Philosophical Inquiry*. Oxford: OUP.
- Kenny, A. (1975) *Will, Freedom, and Power*. Oxford: Blackwell.
- Levy, N. (2011) *Hard Luck: How Luck Undermines Free Will and Moral Responsibility*. Oxford: OUP.
- Nelkin, D. K. and Rickless, S. C. (2017) 'Moral Responsibility for Unwitting Omissions: a New Tracing View', in D.K. Nelkin and S.C. Rickless (eds) *The Ethics and Law of Omissions*, pp. 106–30. Oxford: OUP.
- Rosen, G. (2004) 'Skepticism about Moral Responsibility', *Philosophical Perspectives*, 18: 295–313.
- Sarch, A. (2022) 'Should Criminal Law Mirror Moral Blameworthiness or Criminal Culpability? A Reply to Husak', *Law and Philosophy*, 41: 305–28.
- Sher, G. (2009) *'Who Knew?' Responsibility without Awareness*. Oxford: OUP.
- Sinnott-Armstrong, W. (1984) "'Ought' Conversationally Implies 'Can'", *Philosophical Review*, 93: 249–61.
- Way, J. and Whiting, D. (2016) 'Reasons and Guidance (Or, Surprise Parties and Ice Cream)', *Analytic Philosophy*, 57: 214–35.
- White, A. R. (1975) *Modal Thinking*. Oxford: Blackwell.
- Whittle, A. (2010) 'Dispositional Abilities', *Philosophers' Imprint*, 10/12: 1–23.