

A CAPACITARIAN ACCOUNT OF CULPABILITY FOR NEGLIGENCE

FERNANDO RUDY-HILLER

<https://orcid.org/0000-0002-7977-1216>

Institute of Philosophical Research

National Autonomous University of Mexico

Department of Philosophy

Mexico City

Mexico

ferudy@alumni.stanford.edu

Article info

CDD: 170

Received: 29.09.2021; Revised: 14.01.2022; Accepted: 24.01.2022

<https://doi.org/10.1590/0100-6045.2022.V45N2.FH>

Keywords

Culpability

Negligence

Capacities

Responsibility

Abstract: Ascribing moral and legal responsibility for negligent actions and omissions has always been deeply contested because it seems to be in tension with the natural intuition that responsibility requires control. In this paper I show that we can accommodate culpability for negligence within a control-based account of responsibility if we adopt a “capacitarian” view of control, according to which agents have responsibility-relevant control whenever they have the requisite abilities and opportunity to bring about the morally desired outcome. After explaining the structure of negligent wrongdoing and motivating this conception of control, I show how it can be successfully employed to account for the culpability of negligent agents and to rebut several important arguments against the idea that negligence can be culpable in

the first place. I also explain in what respects my proposal is superior to other capacitarian views found in the literature.

1. Introduction

Ascribing moral and legal responsibility for negligent actions and omissions has always been a deeply contested practice, since it clashes with the very natural intuition that an agent is responsible for all and only those bits of conduct that were voluntarily undertaken and over which they exercised some measure of control (Zimmerman 1986; King 2009; Alexander *et al.* 2009; Moore and Hurd 2011). Since negligent acts and omissions aren't performed "on purpose" but accidentally or at least unwittingly, it follows that attributing responsibility for them is in direct tension with the aforementioned intuition. Alternatively, we can say that, since our moral and legal responsibility practices usually track things like *mens rea*, ill will, malicious intent, etc., on the part of wrongdoers, blaming negligent agents seems to occupy an unstable position within those practices. At the same time, however, negligence is an important concept that moral agents routinely employ to navigate their social world (Amaya 2022) and thus cannot be simply dismissed as an anomaly or an aberration.

Recently, the most powerful challenge to the idea that agents can be morally responsible—and, on the widely held assumption that criminal responsibility is at least partially grounded on moral responsibility, also criminally responsible—for negligence comes from the widely discussed set of arguments put forward by Zimmerman (1997), Rosen (2004), and Levy (2011), according to which unwitting wrongdoing can be culpable only if at its origin lies an episode of clear-eyed wrongdoing. In terms of a slogan, volitionalists—as I have called them elsewhere (Rudy-Hiller

2017)—maintain that culpable wrongdoing can't be unwitting all the way down. In this paper I will sketch and defend a “capacitarian” account of culpability for negligence the main purpose of which is to offer a plausible rebut to the volitionalists' credo, showing how and why moral culpability can indeed be unwitting all the way down and thus providing a secure basis from which to held negligent wrongdoers morally responsible. It's only fair to admit that capacitarian defenses of culpability for negligence have been popular at least since Hart (1968), with other recent proposals having been put forward by Clarke (2014), Stark (2016), Murray (2017), Rudy-Hiller (2017), Vargas (2020), among others. So what's the paper's contribution? I see it as being modest and important at once. It is modest since I won't offer here an entirely new defense of capacitarianism but will mostly rely on mine and other previous defenses of it to show how it is that it can be fruitfully employed to take the edge off several skeptical arguments (including, but not limited, to the volitionalists' arguments mentioned above) against culpability for negligence—whence its importance. This explains why I won't discuss in much detail other capacitarian proposals (with the exception of Clarke's, since it's his that my proposal bears the strongest resemblance): I am not so much concerned here with defending the superiority of my particular strand of capacitarianism, but rather with showing the viability of the capacitarian paradigm in the face of several important recent arguments against it. At the same time, I'll take the time to offer a detailed and (to my mind) much needed clarification of the structure of negligent wrongdoing.

I will offer an account according to which responsibility for negligence is, as volitionalists claim, indeed derivative; the twist, however, is that on my account it doesn't derive from an episode of clear-eyed wrongdoing but from the *unwitting* omission to pay heed to a duty of care the agent was subject

to. Thus far I'm in agreement with Clarke's (2014) account of negligence, but, unlike him, I will argue that this unwitting omission is itself culpable, thus avoiding the implausible view that culpable negligence can be grounded on a non-culpable omission to comply with a duty of care. I will explain how it is that one can be directly blameworthy for unwitting omissions to comply with this sort of duties and also why this feature of my account allows it to successfully block the volitionalists' skeptical argument against the culpability of negligence.

This is the plan of the paper. In section 2 I present an account of the structure of negligent actions and omissions. I argue that the latter are a species of actions and omissions done from morally relevant ignorance and that they are characterized by the agent's prior unwitting omission to comply with a duty of care. I also take the opportunity to explain what duties of care are and why they are central to a correct characterization of negligence. In section 3 I explain why cases of negligence constitute a difficult challenge for control-based account of responsible agency of the sort I defend here. In section 4 I show that the capacitarian account of control I have developed elsewhere (Rudy-Hiller 2017) has the resources for explaining how it is that an agent can be directly responsible for an unwitting omission to comply with a duty of care and derivatively responsible for the resultant negligent action or omission. I also explain there how it is that my account avoids the vicious regress that haunts the volitionalist proposal. In section 5 I engage with another two important arguments that purport to show that there is no culpability for negligence or that such culpability is much rarer than usually assumed, and show that my capacitarian account can rebut each of them. Finally, in section 6 I close with a brief summary.

2. Negligence

2.1 Defining negligence

In this paper I will work with the following definition of negligence:

Negligence. Agent S 's action A or omission O is negligent iff: i) A or O poses an unjustifiable risk of harm to potential victim V or in fact unjustifiably harms V ; ii) S isn't aware that A or O poses such a risk or that it in fact provokes such harm; and iii) the potential or actual harmfulness of A or O is the product of S 's unwitting violation of a duty of care.

According to this definition, negligence is a species of action done from morally relevant ignorance: it is characterized by an action or omission that *unwittingly* poses an unjustifiable risk of harm to others as a result of the agent's *unwitting* breach of a duty of care. Thus, the negligent agent is doubly ignorant: they are unaware that they have flouted a duty of care and they are also unaware that by doing so their action or omission poses an unjustifiable risk of harm to others.

In defining negligence in this way, I side with those theorists who assume that lack of awareness concerning wrongdoing (more specifically, risk of harm) is a distinctive element in negligence—one that distinguishes negligence from recklessness.¹ As standardly conceived, recklessness involves an agent who “consciously disregards a substantial

¹ See for example Sverdlik (1993), King (2009), Alexander *et al.* (2009: Ch. 3).

and unjustifiable risk”² concerning the potential harmfulness of their conduct. In other words, a reckless agent, unlike a negligent one, *is* aware that their action or omission risks harming another person but decides to proceed with it nonetheless. However, despite this apparent clear-cut demarcation between negligence and recklessness, several legal theorists contend that in fact the boundary is a fuzzy one, in large part because, according to them, the dichotomy awareness/unawareness of risk marks opposite extremes in a continuum rather than two discrete states.³ Moreover, some of them claim that there is room for the concept of “advertent negligence” (Moore and Hurd 2011: 149), which would be distinguished from recklessness only by the lesser “grossness” of the unreasonable risk involved. In what follows I will ignore this alleged problem about the fuzzy boundary between negligence and recklessness because my main aim here is to provide an account of the culpability of the kind of wrongdoing characterized by the sort of compounded ignorance described above, regardless of what name we decide to give it in the end.

2.2 Unpacking the definition

I will now briefly unpack the key terms in the above definition of negligence.

² Model Penal Code, Sect. 2.02.

³ See for example Moore and Hurd (2011) and Husak (2011).

2.2.1 *Unjustifiable risk of harm*

I take the expression “unjustifiable risk of harm” from the definition of negligence offered in the Model Penal Code (sect. 2.02). Of course, we have to be careful when transposing the legal conception of negligence to the moral realm, partly because it is standardly assumed among legal theorists that the harm at stake in the legal definition of negligence is restricted to physical injury,⁴ whereas in the moral realm the harm in question can include things such as hurt feelings and disappointed expectations, which are beyond the reach of legal liability. Still, the clause “unjustifiable risk of harm” is relevant in the treatment of moral responsibility for negligence given that, if the agent has a valid justification for a breach a duty of care and consequently for the creation of a risk of harm (or for the actual harm produced), then their conduct isn’t negligent, e.g., if John fails to stop at a stop sign in a residential neighborhood because he is taking a very sick person to the hospital.

2.2.2 *Lack of awareness*

The above definition of negligence introduces unawareness at two different points: first, the agent is unaware that their action or omission creates a risk of harm; second, the risk is the product of a breach of a duty of care concerning which the agent was also unaware (otherwise we have a case of recklessness). As it is standardly conceived by volitionalists (most explicitly by Zimmerman 1997 and Rosen 2004), unawareness in this context must be understood as lack of *occurrent* true belief. Thus, according to

⁴ See for example McBride (2004: 437, n.68).

my definition the negligent agent lacks an occurrent true belief concerning both their having breached a duty of care and their consequently performing an action or omission that poses an unjustifiable risk of harm to others.

Douglas Husak (2011) has argued that we shouldn't understand unawareness in this way, for there are cases in which an agent is occurrently unaware of some risk they have created and yet they do have a belief (albeit a dispositional one) concerning such risk, such as a mother who leaves her baby unattended in the bathtub in order to answer the phone. As the conversation lingers, the mother is occurrently unaware of the risk she has created but, according to Husak, we can't say that she lacks altogether a belief concerning the whereabouts of her baby. Rather, given a suitable prompt, her dispositional belief would easily resurface to consciousness. In this case, Husak claims, it would be implausible to suggest that she has acquired a new belief; rather, what happened was that she became aware of something she believed all along. Husak's conclusion is that, since the mother did have a belief about where her baby was while she was talking on the phone, we should classify her as reckless rather than negligent. He then generalizes and suggests that most, if not all, cases of unwitting risk-creation prompted by a failure of memory should be classified as cases of recklessness, since in those cases agents do have dispositional beliefs about the riskiness of their actions.

This is an important point, but here I will sidestep it for the following reason: given that one of my main goals in this paper is to rebut the volitionalists' challenge against culpable negligence, it is dialectically more fruitful to engage with their own definition of unawareness as lack of occurrent true belief instead of adopting a different one, since if I manage to undermine their argument *even* employing their own conception of unawareness then their view as a whole will lose much more credibility than if I adopted, for instance,

Husak's dispositionalist view. Thus, in what follows unawareness should always be understood as lack of occurrent true belief.

2.2.3 *Duties of care*

According to the definition given above, for an action or omission to be negligent its harmfulness (potential or actual) must be the result of the agent's unwitting violation of a *duty of care*. This is a crucial point for the following reason: it can be the case that an agent performs an action that poses an unjustifiable risk of harm (or actually produces harm) while they are unaware of the harm they are risking (or causing) and yet they can fail to act negligently, provided that there was no duty of care they failed to comply with. For example, suppose that Mary plugs in the coffee maker in the office lounge and that, as a result of a short circuit, she unwittingly brings about a complete blackout resulting in several important electronic files being lost.⁵ Suppose also that the potential failure of the coffee maker was undetectable by a layperson. In this case, even though Mary performed an action that inadvertently brought about unjustifiable harm, she clearly isn't negligent. The reason is simply that she didn't fail to comply with any duty of care that applied to her; for, in the absence of any obvious sign of a potential short circuit (e.g., a stripped wire), it would be absurd to suggest that there is a standing duty to perform extensive inspection procedures before plugging in electronic appliances. So another crucial mark of negligence is the agent's prior failure to comply with a duty of care; this is why negligence is intimately associated with the idea of carelessness.

⁵ Nottelmann (2007: 177-8) presents a similar case.

Now, despite its central importance for understanding negligence, the concept of a duty of care is extremely elusive. It is in the legal literature on negligence where we can find some insight.⁶ To begin with, a duty of care can be defined as “a duty not to harm others by faulty conduct” (Howarth 2006: 450). In turn, faulty conduct is usually understood as conduct that is unreasonable risky, that is, the kind of conduct that fails to accord with what is usually called in the law “the reasonable person standard.” A particular conduct is deemed unreasonable risky when it involves the agent’s failure to take precautionary measures against harm that a reasonable person would have taken. So, for example, if Joanne installs a trampoline in her backyard without a protective mesh, her conduct is unreasonable risky given that a reasonable person (or a reasonably prudent person) wouldn’t have omitted the protective mesh.

There are several debates about the nature of duties of care in the legal literature that, although interesting in themselves, aren’t of direct relevance to the topic of moral responsibility for negligence and so I will set them aside here.⁷ There is, however, an important distinction drawn in legal contexts that *is* potentially significant for the present discussion: the distinction between negative and positive duties of care. McBride (2004: 433) presents this distinction as follows:

As it happens, almost all of the duties of care in negligence that we are subject to are negative in content. The law will only find that A owes B a positive duty of care in negligence—a duty to do something positive for B—if there exists a “special

⁶ See for example Goldberg (2000), McBride (2004), and Howarth (2006).

⁷ See McBride (2004) and Howarth (2006) for a good summary.

relationship” between A and B or if A has “assumed a responsibility” to B.

The distinction McBride is gesturing at seems to be this: while negative duties of care demand that we take steps to avoid the risk of harming some unspecified others, positive duties of care demand that we take active measures to promote the wellbeing of a specific person. It is easy to see why positive duties of care arise only in contexts where there is a “special relationship” between two people: whereas as a driver I have the negative duty of care to take measures to avoid harming pedestrians and other drivers regardless of whether I’m acquainted with them or not, as a parent my duties of care extend well beyond simply avoiding the risk of harming my son; they require in addition taking active steps to promote, for example, his cognitive development. Consequently, I might properly be deemed negligent if, however inadvertently, I fail significantly in the latter task. By contrast, if I inadvertently fail to succor a fellow driver who is out of gas I’m not negligent, since I wouldn’t have breached any duty of care owed to them.⁸

The distinction between positive and negative duties of care can be fruitfully transposed from the legal to the moral realm. In particular, it can explain why moral blameworthiness for negligent actions and omissions is frequently grounded on social roles. The idea is that accusations of negligent wrongdoing in the moral realm often reflect the expectation that the person should have been aware of their obligation to discharge a (positive) duty

⁸ Which doesn’t imply that I’m completely off the hook for my failure to perceive the other driver’s needs. The point is simply that, whatever the nature of my moral failure, it certainly isn’t negligence as I conceive it. Perhaps it reflects a failure of character or lack of virtue.

of care given the special nature of their relationship with another—a relationship that can be grounded on either a social or an occupational role. Of course, it is a further task to specify which relationships give rise to positive duties of care and what it is that those duties exactly demand in each case. The point I wish to make here is simply that the distinction between negative and positive duties of care is useful for realizing the normative importance of social practices and roles for a proper understanding of moral blameworthiness in cases of negligence.

2.3 The structure of negligence

Now we have before us all the elements we need to grasp in more detail the structure of negligence. I agree with Joseph Raz (2011: 265) when he claims that “Negligence involves derivative responsibility: People are responsible for negligent harming because they are responsible for breach of a duty of care.”⁹ Thus, in order to provide a satisfactory account of

⁹ However, I disagree with Raz about the proper account of *what* makes people responsible for unwitting violations of duties of care. In Raz’s view, what he calls the Control Principle is simply unable to explain direct responsibility for unwitting omissions (it is noteworthy, however, that he understands control in volitionalist terms; see Raz 2011: 257, n.7). He proposes instead his “Rational Functioning Principle”, according to which “conduct for which we are (non-derivatively) responsible is conduct that is the result of the functioning, successful or failed, of our powers of rational agency” (2011: 268). And he applies this principle to cases of unwitting omissions: “my suggestion is, we are responsible for omissions due to the malfunctioning of our powers of rational agency” (2011: 267). I think this proposal is inadequate: if an agent’s power to remember some morally relevant consideration malfunctions due to, say, a stroke, it would be wrong to insist that they are nevertheless responsible for the resulting omission. (It’s

culpability for negligence we definitely have to resort to *tracing*, i.e., we need to trace the agent's culpability for their negligent action or omission back to their unwitting violation of a duty of care.¹⁰ However, unlike the volitionalist's tracing

true that Raz says that if the agent in question is mentally disabled then they aren't responsible; being disabled, however, is different from suffering a stroke, since it's different to *have* a mental power that malfunctions than *not having* that power at all—as in the case of the disabled person—. Thus, Raz's view, taken literally, does commit him to the problematic consequence mentioned before.) Something similar occurs if the explanation of the malfunctioning appeals to a situational element (an extraordinary level of noise, for instance) that deprives one of the fair chance of exercising one's abilities. I think that both failures of Raz's account are easily corrected if we switch to the capacitarian account I introduce below, according to which the possession of the relevant abilities and the fair opportunity to exercise them are necessary and sufficient conditions for agents to be responsible for their unwitting omissions and subsequent negligent wrongdoing.

¹⁰ An anonymous referee has objected to the view that tracing is necessary for explaining culpability for negligence by appealing to cases in which “the unwitting violation immediately or constitutively yields the action that creates the unjustifiable risk.” For instance (in a variant of a case to be discussed below), imagine that a person (call him Jim) working in a car rental business forgets to check the brakes of particular car and later on a customer ends up driving the car in suboptimal conditions. I agree with the referee that Jim is culpable of negligence, but disagree with them that no tracing is needed to explain his culpability. This is because Jim negligently *omitted* to check the car's brakes *as a consequence* of having forgotten to pay heed of a duty of care he was subject to. In a different kind of case presented by the referee, Jim is driving his car when he suddenly realizes that he failed to take it to the mechanic to have its brakes checked on time. Again, Jim is culpable of negligence but, again, tracing is needed to explain why. Here, Jim negligently *drove* a car with poor functioning brakes *as a consequence* of having forgotten to pay heed of a duty of care he was

account of unwitting wrongdoing developed by Zimmerman (1997) and Rosen (2004), on my view as well as Raz's (2011: 266) the breach in question is an unwitting omission—the unwitting omission to comply with a duty of care.¹¹ Therefore, if we are to vindicate commonsense regarding culpability for negligence (which is exactly what the volitionalist does *not* try to do), we must explain how it is that agents can be *directly* responsible for their unwitting omissions to comply with duties of care. This is the central task for an account of responsibility for negligence and I'll undertake it in section 4.

Taking our cue from Raz's statement, we can present the basic structure of a case of negligent wrongdoing in this way:

- i) At t_1 agent S unwittingly omits to comply with a duty of care they are subject to;
- ii) at t_2 , and as result of the prior unwitting omission, the agent performs an action that harms, or risks harming, someone; or omits to perform an action, and such omission harms, or risks harming, someone;
- iii) the agent is unaware at t_2 that the action they perform, or the omission to perform one, harms, or risks harming, someone.

subject to. Therefore, in either of the cases presented by the referee, tracing is indeed required to account for the agent's culpability.

¹¹ Clarke (2014: ch. 7) also emphasizes that negligent actions or omissions are usually preceded by an unwitting omission to comply with a duty of care, though he is wrong in claiming that the latter needn't be culpable for the agent to be responsible for their negligence. I explain why in subsection 4.4 below.

Let me illustrate this account with a couple of examples, one corresponding to negligent action and the other to negligent omission. Suppose that Jim forgets to take his car to its annual checkup and as a result ends up driving it for several weeks with poor functioning brakes. This is a textbook case of negligent action: driving a car in these conditions is negligent because Jim unwittingly performs an action that risks harming other people and the riskiness of his action is a consequence of an unwitting failure to comply with a duty of care he was subject to.

Now think about Susan, a forest ranger. Susan is required to be in a certain specific post during heavy storms in order to be able to alert of incipient wildfires due to lightning bolts. During a particular storm, however, she forgets about her duty and as a result she omits to alert of an incipient fire. Unfortunately, her omission ends up causing substantial harm.¹² This is a textbook case of negligent omission: Susan's unwitting omission to alert about the fire is negligent because the omission causes harm and yet she is ignorant both of her omitting to do something she was required to do and of the harmfulness of such omission, and all this occurs as a result of her unwitting failure to pay heed to a duty of care she was subject to.

3. The problem with negligence

The best way to understand why responsibility for negligence appears to be deeply problematic is to adopt for a moment the volitionalist's perspective. According to the volitionalist's view, an agent has direct control over—and so is directly responsible for—an action or an omission only if

¹² For a discussion about whether omissions can cause anything, see Clarke (2014: 54-8).

they perform the action or the omission in full awareness of what they are doing and its moral valence (Zimmerman 1997). Given this understanding of control, it is only natural that responsibility for negligence seems difficult to make sense of. Manuel Vargas (2020: 402) explains the difficulty in this way:

It may seem unclear how to accommodate our attitudes about negligence in a control-based account. One's neglect of the morally preferable option is, in the ordinary case, non-volitional. We do not intend to fail to consider some salient consideration, we do not plan to ignore demands for due care, and we do not usually seek to forget about shared commitments ... In typical instances of negligence, the relevant consideration simply does not spring to mind.

As Vargas goes on to note, from the volitionalist's perspective it is mysterious why the agent's failure to recognize the relevant considerations and act on them is supposed to be culpability generating rather than exculpatory. The only answer the volitionalist can countenance appeals to tracing the negligent action or omission back to an episode of clear-eyed wrongdoing whereby the agent *decided* to disregard pertinent considerations such as the obligation to pay heed to a duty of care, a view which is obviously in conflict with the standard conception of negligence as involving carelessness rather than maliciousness. Thus, the problem with negligence is this: proceeding under the assumption that some form of control seems to be an indispensable condition on any correct account of responsible agency,¹³

¹³ There are of course some philosophers, usually called "attributionists," who roundly reject the centrality of control for

how can we make sense of the fact that in our ordinary practices we are deeply committed to cases of nonvolitional (or nonvoluntary) culpability?

From the volitionalist's perspective, negligence is particularly puzzling since it is an instance of wrongdoing that is *doubly* nonvolitional or nonvoluntary. First, the negligent agent doesn't disregard a duty of care voluntarily; second, they don't create a risk of harm voluntarily through their action or omission either. I claimed above, following Raz, that responsibility for negligence derives from the agent's *direct* responsibility for the unwitting omission to discharge a duty of care. Thus, the central challenge to unravel the mystery of negligence is to explain how it is that agents can retain direct control in cases of unwitting omissions. It is to addressing this challenge that I now turn.

4. A capacitarian account of culpability for unwitting omissions and negligence

My main contention in what follows will be that a "capacitarian" account of control is able to explain direct responsibility for unwitting omissions to discharge duties of care, as well as derivative responsibility for negligence, in a way that blocks the volitionalist's regress argument. In this section I will present my proposal and in section 5 I will expand on its virtues by discussing other skeptical arguments about the culpability of negligence found in the literature.

responsibility. See for example Scanlon (1998) and Smith (2005). I don't find their arguments compelling, but I won't discuss them in this paper. Rather, I will simply proceed under the assumption that some form of control is indeed essential for responsible agency.

4.1 *Capacitarian control*

Consider the following conception of responsibility-relevant control:

Capacitarian control. An agent *S* has direct control over φ if and only if *S* has the requisite cognitive and executive abilities to make it the case that φ occurs and a fair opportunity to exercise them.

What can be said in favor of such a view? Consider as an analogy the case of a soccer goalkeeper who fails to stop an easy ball of the sort they have easily caught many times before, and suppose also that all the conditions were favorable to their doing so this time as well: they retained all the relevant abilities for stopping the ball and there were no unusual external factors (e.g., a heavy storm) obstructing their correct exercise. In this sort of case, it's only natural for fans and teammates to blame (in a non-moral way, of course) the goalkeeper for their failure and to think of them as directly responsible for their failure to properly exercise the capacities they do possess. And as in sports, so it is in morality: we routinely blame each other for our failures to perform competently with respect to certain basic moral demands—including, importantly, our duties of care—whenever we are deemed “normal” moral agents possessed with the relevant cognitive and executive abilities to comply with those demands and our environment is thought to be favorable, or at least not severely inimical, to our exercising them. This feature of our responsibility practices is the backbone of the capacitarian view of control.¹⁴

¹⁴ I have defended this view of control in detail elsewhere (Rudy-Hiller 2017), focusing on cases of culpable ignorance rather than

Now it's true that, in the previous paragraph, it is assumed that moral responsibility tracks whether or not the agent *has* control and not whether the agent *exercises* control at the time of wrongdoing. And this may seem problematic because volitionalists insist that the latter, rather than the former, is the central notion, which is why they deny that negligent agents—that is, agents whose wrongdoing derives from an unwitting omission—could be morally responsible, since unwitting omissions are paradigmatic cases of lack of exercise of control (Zimmerman 1986). So, is my argument simply begging the question against volitionalists?

I don't think so. Here is why: to being with, although I readily admit that I don't have a knock-down argument against volitionalism (since volitionalists can always insist that the capacitarian conception of control fails to satisfy them), it's crucial to note here that my arguments aren't primarily addressed to volitionalists themselves but rather to those who are worried about the skeptical upshots of their position and would like to explore viable alternatives. This doesn't entail, however, that there's nothing we can say *against* volitionalism in order to break a potential "dialectical stalemate" between the capacitarian and the volitionalist conceptions of control. What we can do is to *raise the cost* to the volitionalist for adopting a revisionist stance toward our responsibility practices and, more specifically, toward the practice of blaming negligent wrongdoers, by showing *how important* for social creatures like us those practices are. In particular, following Raz's (2011) and Murray and Vargas' (2020) proposals, I suggest we should focus here on the interest negligent agents *themselves* have in being held to account for their wrongdoing.

negligence proper. Thus, this paper represents an important extension in the scope of my view.

The idea is this. It isn't only the people surrounding the unwitting wrongdoer who have an interest in making them a more dependable agent partly by blaming them; the agent themselves also have a particular interest in being held accountable. Raz (2011: 268) nicely articulates the interest in question as follows: "In acknowledging our responsibility for these unintentional acts and omissions we affirm our mastery of these abilities, deny that we are disabled in the relevant regards." The idea is that the unwitting wrongdoer is prepared to accept blame because doing so is a way of reaffirming their status as a competent moral agent. Conversely, refusing to accept blame on the grounds that one "simply forgot" or "didn't notice" or "failed to think about" the relevant considerations can rightly be taken by others as a statement that one lacks the relevant abilities, and so as a statement that one is less than fully competent in moral matters. However, given the kind of social creatures we are, projecting and accepting such a degraded moral status have considerable costs, and so it is natural that agents normally have an interest in avoid doing so. Raz's important point is that this interest plays a central role in the agent's readiness to accept blame in the kinds of cases that are of interest here.

There is a related reason why agents would want to accept blame in cases of unwitting wrongdoing—a reason that has to do with avoiding an alienated view of one's own agency. If an agent refuses to accept responsibility for an unwitting omission on the grounds that, for example, the failure to recall the relevant considerations was due to the malfunctioning of a subpersonal mechanism that has nothing to do with themselves, then this can rightly be interpreted as the agent's taking a detached view of their agential capacities—a view that seems incompatible with a healthy unified conception of oneself. Such an agent, if consistent in adopting this view, would be prone to

schizophrenic-sounding statements like: “Oh, don’t praise me for this witty comment; it just popped up in my head as the result of some mechanism inside me!” Similarly, in cases of unwitting wrongdoing such an agent would be prone to externalize blame, that is, to refuse to acknowledge that the relevant failure reflects something important about *them*. Conversely, being prepared to accept blame for unwitting wrongdoing is a sign of one’s taking the said healthy unified view of oneself and of one’s agential capacities.

Once again, this isn’t meant to be a knock-down argument against volitionalism; rather, it’s a pragmatic argument for taking our practice of blaming negligent wrongdoers seriously which aims at making apparent the high costs that accepting the volitionalist view of control would entail.

4.2 Applications

Now let’s put to work the capacitarian conception of control to explain how is that *direct* responsibility for unwitting omissions is possible.¹⁵ My key contention will be that an agent is directly responsible for an unwitting omission to comply with a duty of care whenever is the case that they *could* have complied with it—in a sense of “could” that is made precise by the capacitarian account.

Take Jim’s case again. It is time for his car’s annual checkup. Given this fact, plus Jim’s projected actions that involve routinely driving his car and the possible harmful consequences of doing so if the car is in bad shape, he ought to take it to inspection. He is under a duty of care to do so.

¹⁵ This is the point where my view bears the strongest resemblance to Clarke’s (2014, 2017a, b). In 4.4 below I explain where the two views differ and why mine is superior.

Now suppose that he fails to realize this; that is, he fails to recall that it is time for his car's annual checkup and as a consequence fails to realize that he ought to take his car to inspection. As a further consequence, he unwittingly omits to do so. If he goes on to drive his car in its present shape, including its suboptimal brakes, then his action is negligent. I claimed above that he is derivatively responsible for such action—and for some of its consequences—because he is directly responsible for the unwitting omission to take his car to inspection. I will substantiate the latter claim with the help of the capacitarian account of control.

Recall that the capacitarian account's basic idea is that an agent has direct control over an action or omission¹⁶ if and only if they have the requisite cognitive and executive abilities to perform the action in question and nothing decisively interferes with their deployment—which is to say, the agent has a fair opportunity of exercising those abilities.¹⁷

¹⁶ And over cognitive operations that aren't actions such as noticing or remembering relevant considerations. See (Rudy-Hiller 2017) for arguments in favor of this view.

¹⁷ An anonymous referee has worried that my use of the notion of fair opportunity might be “awfully circular” if “fair” just means “when it is fair to blame the agent for not exercising the capacities they have”. This is a (sorry for the pun) fair worry, but I'm not using the notion in the circular way that troubles the referee. Rather, building on Brink and Nelkin's (2013) work, elsewhere (Rudy-Hiller 2020) I have offered an understanding of fair opportunity in terms of *situational aptness*, that is, a situation in which the usual known impediments for a proper exercise of our cognitive and executive abilities are absent or at least their influence is negligible, so that it's *reasonable to expect* the agent to adjust their behavior to certain (moral) demands by exercising those abilities. In turn, the notion of “reasonable expectation” is partly normative and partly descriptive: it's normative because it incorporates moral demands, but it's also descriptive because it

So if Jim had direct control over his omission to take his car to inspection despite the fact that it was an unwitting one, it must have been the case that he had the abilities required to recall that it was time for his car's annual checkup, to recognize the moral significance of this fact and, finally, to act on this recognition and take the car to inspection. And it must also have been the case that there weren't situational factors that made it unusually hard for him to exercise these abilities. Are all of these requisites fulfilled in Jim's case?

A commonsensical view of unexercised capacities and fair opportunity answers this question as follows. If Jim is a normal human being and his situation wasn't significantly out of the ordinary, the answer is affirmative. For, in the absence of evidence to the contrary, we may suppose that Jim is endowed with the average powers to remember significant facts, to foresee some of the most salient probable consequences of his actions and omissions, and to recognize and respond to moral considerations. In the case at hand, remembering that it was time for his car's revision is something he was capable of doing, not only because he had remembered it many times in the past but also because he routinely remembers significant dates related to his duties (e.g., tax day, semi-annual dentist appointments). He was

incorporates empirical information about what kind of factors negatively affect human moral capacities, including (when we have them) *base rates* of certain kinds of behavior drawn from experimental social psychology (such as "situationist" experiments). This appeal to base rates also allows us to say that, rather than being a "yes or no" notion, fair opportunity comes in degrees (and so does blame), and therefore it's more appropriate to talk about the *quality of opportunity* the agent has for behaving properly (Nelkin 2016; Rudy-Hiller 2020). In sum, the notion of fair opportunity as I use it has considerable content that goes far beyond a mere appeal to intuitions about when it's fair to blame agents for not exercising their capacities.

also able to foresee the probable consequences of failing to take his car to inspection, since he regularly deploys the capacity to foresee consequences in a host of different situations. And he was able as well to recognize the moral significance of those probable consequences, namely, that he had a moral reason—grounded on a duty of care—to take his car to inspection. Finally, he was able to actually take his car to inspection, since he has the usual volitional powers to form and execute intentions on the basis of reasons—moral and nonmoral—he has recognized.

Additionally, if there was no feature of Jim's situation that decisively impaired the exercise of these abilities—for instance, if it wasn't the case that he was so overwhelmed by grief that it was extremely difficult for him to summon the relevant thought about his car; he didn't suffer a cerebral aneurism that left him unable to think about the consequences of his actions; he wasn't paralyzed and couldn't have taken the car to the mechanic even if he had tried—then we can say: Jim could have remembered the relevant date, could have foreseen the consequences of not taking his car to inspection, could have recognized his having a duty of care to do so, and could have taken his car to inspection. In sum, Jim could have done otherwise—which is to say, he had the requisite capacities and opportunity to do otherwise. It is this fact which, according to the capacitarian account, makes it appropriate to hold him *directly* blameworthy for the unwitting omission to comply with his duty of care and take his car to inspection and *derivatively* blameworthy for the resulting negligent action of driving his car in suboptimal conditions.¹⁸

¹⁸ It's worth noting that, according to the capacitarian proposal, neither *mens rea*—guilty mind—nor ill will are necessary for blameworthiness. According to self-expression theorists, this is a problematic consequence of capacitarianism and so they are also drawn to a skeptical view about responsibility for negligence (see

A similar explanation can be given of Susan's direct responsibility for her unwitting omission to be at her post during the storm. In the absence of evidence to the contrary, we are entitled to attribute to her the cognitive and executive abilities needed to recall her duty and to act on this recognition. And absent conspicuous situational factors that could excuse her from her failure to exercise those abilities, we can also attribute to her a fair opportunity to exercise them. So we can legitimately say that Susan could have been at her post during the storm despite her unwitting omission to do so; that is, she could have done otherwise, and so she is directly blameworthy for this omission. And she is also derivatively blameworthy for one of its consequences, namely Susan's negligent omission to alert of the incipient wildfire.¹⁹

It is worth mentioning that the relevant sense of "can" at issue in statements like "Jim could have recalled to have his car's brakes checked" is *not* what is sometimes called the metaphysical sense of "can" (Zimmerman 1986). Thus, I am not claiming merely that it's metaphysically possible for Jim to have done so—although of course it is at least that—since this would clearly be insufficient to ground his responsibility

esp. Talbert 2017). For lack of space, I will have to defer a detailed treatment of this view for another occasion.

¹⁹ For an agent to be responsible for the consequences of failing to comply with a duty of care it must be the case that those consequences are of the kind the duty of care in question was meant to prevent. In both Jim's and Susan's cases, I take it as obvious that their respective duties of care were meant to prevent exactly the kind of bad consequences that occurred, namely, Jim's driving his car in suboptimal conditions and Susan's failure to alert of the incipient fire.

for the negligent action of driving a car with faulty brakes.²⁰ Rather, I am claiming that—to use Zimmerman’s term—is “personally possible” for Jim to have complied with his duty of care given that he had all the requisite abilities and the opportunity to have done so. Zimmerman thinks that this sense of “can”—the one involving personal possibility—is restricted to intentional actions and omissions. But, again, this is true only if we assume that the “can” of personal possibility is necessarily linked to the volitionalist notion of control. If we reject the latter and adopt the capacitarian notion instead, then the range of what an agent can do—in the sense of personal possibility—is expanded considerably.

4.3 *Blocking the regress*

As we have seen, the capacitarian account just sketched is able to explain responsibility for negligence without needing to invoke an episode of clear-eyed wrongdoing from which the agent’s blameworthiness derives and so we needn’t have to suppose that the negligent agent *decides* to dodge one of their duties of care in order for them to be culpable for their negligence. Consequently, it blocks the regress that, according to volitionalists like Zimmerman, Rosen, and Levy must be pursued until such episode of clear-eyed wrongdoing is found. On my account, the regress from

²⁰ This is because there are many things that are metaphysically possible for us to do and yet for which we aren’t responsible, because we lack either the responsibility-relevant capacities or the fair opportunity to exercise them, or both. For example, it may be metaphysically possible for me to resist the threat “rob the bank or I kill your wife” without its being the case that I have a fair opportunity to do otherwise than rob the bank, in which case I am not blameworthy for doing so.

the negligent action or omission to the culpability basis terminates whenever there are sufficient grounds for deeming that the agent had capacitarian control over the unwitting omission to comply with one of their duties of care. When this occurs, the agent is directly blameworthy for the unwitting omission and derivatively blameworthy for any subsequent negligent action or omission.

A virtue of this account is that it vindicates everyday attributions of responsibility for negligence. I take it as given that, absent some compelling excuse, we would judge that Jim is blameworthy for driving his car in suboptimal conditions and for some of the consequences of his action. Similarly, we would judge that Susan is blameworthy for failing to alert of the incipient fire and for some of the consequences of her omission. My account is able to explain why these judgments are warranted: these agents are blameworthy because they had responsibility-relevant control over the actions they unwittingly failed to perform and that would have led them to avoid their negligent wrongdoing. They could have done otherwise and so they are responsible for their negligence—regardless of the fact that it was committed unwittingly.

4.4 Clarke's account

My account of moral responsibility for negligence bears some resemblance to Clarke's (2014, 2017a, b) proposal, since both of us provide a capacitarian response to the volitionalists' skeptical argument. However, there is a key difference between our accounts that needs to be acknowledged in order to clearly appreciate the distinctiveness and advantages of my view. Clarke (2017a: 245-6; 2017b: 72-3) thinks that agents can only be directly responsible for failures to perform bodily actions or

omissions and can never be directly responsible for cognitive failures like forgetting an important piece of information or failing to recall the need to discharge a specific duty; crucially, however, these are exactly the kind of cognitive failures that lie behind the negligent agent's failure to comply with their duties of care, so in denying that agents can be blameworthy for these failures, Clarke is committed to the view that culpability for negligence is grounded on something for which the agent isn't culpable, i.e., their failure to comply with a duty of care.²¹

To see why this is implausible, consider the following. Clarke (2017b: 72-3) claims that these failures, though not themselves culpable, are "morally faulty" and that responsibility for the negligent action or omission is, contrary to what I have been arguing, direct rather than derivative. In our examples, Jim would turn out to be directly responsible for his negligent action of driving a car with poor-functioning brakes and Susan directly responsible for her negligent omission of not alerting of the incipient wildfire, though in both cases their negligence would be grounded in a moral failure for which they aren't responsible—the failure to discharge their respective duties of care. But now the problem is this: when it comes to moral faults which aren't themselves blameworthy—like characterological failures and vices—the appropriate reaction is, as Clarke (2017a: 245) himself admits,

²¹ An anonymous referee has objected that, since the cognitive failures that prompt the unwitting violation of a duty of care can be distinguished from the violation itself, we don't have a necessary "transfer" of blame from one to the other. I agree, but my argument against Clarke's view doesn't appeal to this problematic kind of transferring. Rather, as I explain in the text below, my argument appeals to the untoward consequences of adopting a "characterological" view of those failures.

admonition (i.e., “you shouldn’t be that way”) rather than outright moral blame and its accompanying reactive attitudes like resentment and indignation. But if this were right in the case of negligence, it would turn out that when we hold negligent wrongdoers accountable what we are actually doing is something like admonishing them for having those reprehensible character traits that led them to disregard their duties of care in the first place rather than blaming them for what they negligently did or omitted to do.²² Thus, far from being a vindication of moral responsibility for negligence, Clarke’s proposal ends up resembling very closely Moore and Hurd’s (2011) *skeptical* view of the culpability of negligence (to be discussed in the next section) according to which blame in the case of negligence actually comes to no more than faulting agents for being “stupid, clumsy, selfish and weak”.

In order to avoid this “characterological” view of negligence, what we should do is to pursue the capacitarian strategy to its ultimate consequences and claim, as I have done above, that agents can be *directly* blameworthy for their failures to discharge duties of care whenever is the case that they have the capacities and opportunities needed for having paid heed to them and even when those failures involve non-

²² An anonymous referee wondered “why not just think that we admonish for the underlying cognitive failures but at the same time we blame or punish for the negligent risk creation. That we might admonish for the one does not mean that we cannot blame for the other.” In response, notice that, in the kind of view the referee sketches, blame is bound to be seen as unfair because it would be grounded on a character fault for which the agent is very likely *not* responsible such as, for instance, obliviousness or stupidity. But, as I point out in the text, this is exactly the kind of argument Moore and Hurd (2011) present *against* the culpability of negligence. Thus, the referee’s proposal would be problematic as a *defense* of that culpability.

actional omissions like failing to remember a crucial date (as in Jim's case), and *derivatively* blameworthy for the negligent actions or omissions derived from that failure. In this way, blame, rather than admonition, gets firmly rooted at the origin of the agent's unwitting wrongdoing and the characterological view of negligence is thus avoided.²³

5. Two other arguments against the culpability of negligence

In this section I will review two important recent arguments, offered by Moore and Hurd (2011) and King (2009), that, independently of volitionalism, purport to establish that there is no genuine moral responsibility for negligence. I will argue that the capacitarian account presented above is able to successfully rebut each of them, thus strengthening its credentials as a proposal worth taking seriously.

²³ Nelkin and Rickless' (2017) account of responsibility for unwitting omissions has important points in common with Clarke's (2017a, b) but, unlike Clarke, they employ tracing to explain culpability for unwitting wrongdoing. However, and contrary to my own view, in Nelkin and Rickless' account the requisite anchor for direct responsibility necessarily involves awareness of the fact that one could do something to prevent one's future negligent wrongdoing, for instance writing reminders to ensure one doesn't forget to take one's car to its annual check-up. This feature of their view renders it unfit for cases of negligence as I have defined them, because in these cases the agent is unaware of the risk of wrongdoing not only at the time of action but also at every relevant time in the past.

5.1 Moore and Hurd's attack on the culpability of unexercised capacities

In their provocative paper “Punishing the Awkward, the Stupid, the Weak, and the Selfish: The Culpability of Negligence”, Michael Moore and Heidi Hurd (2011) aim to show that there is no moral—and consequently no criminal—responsibility for negligence. Moore and Hurd explore at length five different explanations of the culpability of negligence and conclude that none of them provide a justification for blaming and punishing the inadvertent creation of unreasonable risk. One of the possible explanations they explore and dismiss is precisely the one I have defended here—that the negligent agent is culpable because they had, but failed to exercise, the relevant capacities to advert to the risk of harm and do something to prevent it. Though they focus exclusively on Hart’s (1968) proposal, they adventure a far-reaching conclusion: “there simply is no defense to make of the blameworthiness of negligence in terms of there being a culpability of unexercised capacity” (Moore and Hurd 2011: 165). I will show that this conclusion is based on an inadequate conception of how unexercised capacities ground blameworthiness.

To begin with, Moore and Hurd (2011: 158-60) assume that unexercised capacities must necessarily be understood in terms of the classic conditional analysis of ability. According to the conditional analysis, to say that agent S had the ability to φ in a case where they didn’t actually φ means something like: “S would have φ -ed if they had chosen (willed, desired, etc.) to φ ”. Generalizing from this, Moore and Hurd claim that any statement of the form “S could have φ -ed” (where “ φ ” doesn’t necessarily stand for an action verb; “ φ ” can be replaced by verbs like “advert”, “notice”,

“remember”, etc.) requires the truth of a statement of the form “S would have φ -ed if C”, where “C” stands for the condition that must be satisfied for it to be true that S would have φ -ed. In the case of abilities to perform intentional actions like running a four-minute mile, it is assumed that the relevant condition is the agent’s choice to exercise the relevant ability. However, when the ability in question is the ability to advert to the risk of harm, as it is in cases of negligence, the relevant condition can’t be the agent’s choice since it isn’t true that the agent would have adverted to the risk of harm had they chosen to do so (adverting can’t be done at will). So in order to make sense of the unexercised ability to advert to the risk of harm, we need to find out what occupies the “C” place in the relevant counterfactual. Moreover, whatever occupies this place must be a “morally plausible desert basis for blame” (Moore and Hurd 2011: 160), since the whole point of the analysis is to find out what it is that makes negligent agents blameworthy.

In their view, the relevant condition in counterfactuals about the unexercised capacity to advert to the risk of harm appeals to “those characteristics of the inadvertent agent’s practical rationality that explain his inadvertence” (Moore and Hurd 2011: 166). They divide these characteristics into four main groups: motivational, cognitive, conative, and related to motor skills. They then claim that statements of the form “S could have adverted to risk X” must be interpreted in the model of any of the following four statements (Moore and Hurd 2011: 164-5):

- a) S would have adverted to X if he wasn’t so selfishly indifferent to the welfare of other people (motivational failure);

- b) S would have adverted to X if he wasn't so stupid as to fail to infer the risk of harm from his proposed action (cognitive failure);
- c) S would have adverted to X if he wasn't so weak-willed as to being easily distracted by trivial incentives (conative failure);
- d) S would have adverted to X if he wasn't so clumsy²⁴ (motor skill failure).

Moore and Hurd claim that these four kinds of defects of practical rationality or characterological shortcomings not only constitute the relevant condition in counterfactuals about the unexercised capacity to advert to the risk of harm but also constitute a *prima facie* morally plausible desert basis for blame. But, they contend, there are two problems with this suggestion: first, if these defects are what really ground the culpability of negligence, then unexercised capacities do no real work to explain the latter; and second, if this were correct, then when we blame negligent wrongdoers what we would actually be doing would be faulting them for having bad characters and, of course, having a bad character isn't, in and of itself, an appropriate basis for holding agents responsible and much less for subjecting them to criminal liability (Moore and Hurd 2011: 164).

The capacitarian account of culpability for negligence I presented above has the resources for showing where it is that Moore and Hurd's argument goes astray. Begin by noting that my account doesn't rely on the conditional

²⁴ Their idea here is that some episodes of inadvertence are caused by the agent's poorly executed actions. For instance, a motorist may fail to notice a "Children playing" sign because they dropped their cigarette just as they were passing by it.

analysis of ability. Susan Wolf, a prominent capacitarian, wrote the following to distinguish her approach from the traditional one: “Although I can offer no reductive analysis of ‘ability’ ... I can give a *characterization* of what is involved in attributing an ability to someone” (Wolf 1990: 101). This is precisely what I did above: I offered a characterization of the attribution of unexercised abilities in cases of unwitting omissions without aiming to provide a reductive analysis of what it is to have an unexercised ability. But if the capacitarian account, properly understood, needn’t rely on the traditional conditional analysis of ability, then Moore and Hurd’s argument has no force whatsoever. For then the capacitarian can resist their contention that statements of the form “S could have adverted to risk X” must necessarily be understood in terms of counterfactuals like “S would have adverted to risk X if S wasn’t so selfishly indifferent to the welfare of others.” Rather, it is open to the capacitarian to insist that the former statement can be true given conditions precisely as they were²⁵—that is, holding fixed all the relevant features of the agent’s practical rationality. This means we needn’t suppose that the truth of such statements presupposes substantive changes in who the agent is; all that is needed is that the “modal profile” of the agent’s capacity to advert to the risk of harm is such that we can legitimately say that they could have adverted to it when they failed to do so, given their abilities and the opportunities that were present to them. So while it is true that one way to determine such profile is by appealing to counterfactuals and possible worlds, it is false that, as Moore and Hurd (2011: 165) claim,

²⁵ I take the phrase “given conditions precisely as they were” from Wolf (1990: 100), who is discussing Austin’s claim that not all ability statements are conditional.

all such counterfactuals “presuppose some changes in who [the agent] is.”²⁶

If the truth of statements like “S could have adverted to risk X” can be ascertained without having to resort to a conditional counterfactual of the sort employed in the classic condition analysis of ability, we needn’t embark on Moore and Hurd’s quest for some morally appropriate desert basis for blame in cases of negligence beyond the fact that the agent in question did have the relevant capacities and had, in addition, a fair chance to exercise them. So unexercised capacities can ground blameworthiness for negligence directly, without having to appeal to the agent’s faulty character. This means that unexercised capacity talk does have an ineliminable role to play in explaining responsibility for negligence and also that there is such a thing as the culpability of unexercised capacities.

Now Moore and Hurd could insist that even conceding the irrelevance of a conditional counterfactual for accounting for the culpability of negligence, the crucial point is that a feature of the agent as they actually are—for instance, their selfishness—is what ultimately explains that culpability, since it is this feature what explains why the agent failed to exercise the relevant capacity. I don’t think this is correct. To see why, consider this scenario. An agent negligently causes harm, or negligently risks causing harm, and yet fails to evince any of the four failures of practical rationality Moore and Hurd listed. That is, the agent’s inattentiveness isn’t explained by selfishness, stupidity,

²⁶ See Vargas (2013: 215, n. 22): “[I]n characterizing the relevant capacities in this way [by appealing to possible worlds], the capacities are not capacities of other agents, or idealized agents, or counterfactual agents. Rather, what responsibility-relevant capacities agents have *in the actual world* is settled by an idealized set of counterfactuals.”

weakness of will, or clumsiness. In other words, the agent's failure to advert to the risk of harm was a brutally explainable failure to exercise the relevant capacity (McGeer and Pettit 2015). Even if this is so, I think that it would be perfectly appropriate to hold them responsible for the harm (caused or risked) provided that we have good grounds for thinking that they were capable of doing otherwise—that is, that they had the requisite capacities and fair opportunity for doing otherwise.

Recall the goalkeeper's example discussed before. The goalkeeper can be blameworthy (in a sports sense rather than in a moral sense) for their failure to stop an easy ball even if this failure can't be explained by a characterological failure on their part like indifference or chronic lapses of concentration. On the contrary, the goalkeeper can be a model of professionalism and commitment and yet, if they had the requisite capacities and opportunity for catching the ball, they can rightly be blamed for the slip that led to their team's defeat. Of course, if the teammates and the fans discovered that the goalkeeper was lacking in commitment to the team, that would increase their resentment towards them; but, as I have just argued, the goalkeeper would still be blameworthy for the defeat even if no such characterological failure can be found. Thus, in sports as in morality, there is culpability of unexercised capacities independently of those defects of practical rationality that, in some cases, explain the failure to exercise the relevant capacity and that may enhance the person's blameworthiness for the harm caused or risked but aren't necessary for it.

5.2 King and the ad hocness charge

My capacitarian account of negligence also has the resources for countering Matt King's (2009) widely discussed

argument against the culpability of negligence. King's main contention is that all possible explanations of the latter are bound to be objectionable *ad hoc*:

if we are unable to give a general explanation of responsibility that can account for cases involving negligence, either negligence requires an exceptional *ad hoc* explanation, or else we ought to reject the claim that negligent agents are responsible for the harms they bring about (2009: 582).

King is convinced that no general explanation of the culpability of negligence can be had because, on the one hand, what he calls "paradigmatic cases" of responsibility involve conscious mental states (for example, about the action's moral valence or about its consequences) and, on the other, negligence is *defined* by the lack of such states (King 2009: 581). So it follows that whatever explanation we offer in cases of negligence it would have to be very different from the one we offer in paradigmatic cases. But such an explanation would "ignore the fact that in both cases [paradigmatic and non-paradigmatic] the agent involved is supposed to be responsible for x" (King 2009: 587), which implies that there must be a certain unity in our theory—the kind of unity which would be lacking from the envisioned account.

Once again, the capacitarian account of control has the resources to respond to this skeptical argument about the culpability of negligence. The response comes in two stages. First, the capacitarian theorist denies that paradigmatic cases of responsibility are characterized by actual awareness concerning wrongdoing. Rather, they insist that responsibility requires the possession of certain cognitive and executive abilities and a propitious situation for their

deployment. If the agent and their situation meet these conditions, then the former is responsible; otherwise, they aren't—unless, of course, the agent are themselves responsible for bringing about the deterioration of their capacities or for making the situation unfavorable to their deployment.

Now King and other volitionalists could object in the following terms: if the paradigm of responsibility isn't an agent who performs an action intentionally and in full awareness of its moral valence, then why do we assign more blame in an instance of advertent and intentional wrongdoing than in an instance of negligence? As a response, note that acknowledging this fact about our attributions of blame doesn't require adopting the volitionalist picture. Rather, we can stick to the capacitarian account and claim (with the Strawsonians) that attributions of blame are sensitive to the perceived degree of ill will on the offender's part. Then we can note that the perceived degree of ill will usually co-varies with the degree of awareness we presume the agent had concerning the wrongness of their action. Given this fact, it isn't surprising that we blame more harshly a witting than an unwitting wrongdoer, but this shouldn't be taken to imply that the former satisfies paradigmatic conditions of responsible agency that the latter doesn't. If both had the relevant capacities and opportunity, both meet the required conditions for responsibility in equal measure.

The second stage of my response to King consists in arguing that the capacitarian account does offer a unified explanation of responsibility encompassing cases of intentional wrongdoing and negligence. According to this explanation, the unifying element consists in the notion of capacitarian control: both witting and unwitting wrongdoers are responsible if and only if they have the kind of control that characterizes responsible agency. As we have seen, the

capacitarian understands control in terms of the possession of the requisite abilities and the fair chance to exercise them. Thus, this explanation provides a unified account of responsibility in terms of the possession of certain capacities and certain features of the agent's situation that together amount to the responsibility-relevant notion of control. So once we deny the volitionalist conception of paradigmatic cases of responsibility, and once we offer an alternative picture of responsible agency—a picture that applies both to witting wrongdoers and negligent ones—we have successfully provided an account of responsibility for negligence that isn't *ad hoc* at all.

6. Conclusion

In this paper I presented and defended a capacitarian account of responsibility for unwitting omissions and negligence. After offering a conception of the structure of negligent wrongdoing in terms of derivative responsibility that originates from the unwitting breach of a duty of care, I argued that the central problem for explaining culpability for negligence is precisely to explain how an agent can be directly responsible for an unwitting omission given that direct responsibility entails direct control and that it is hard to see how an agent can have direct control over an unwitting omission—especially if one adopts the volitionalist notion of control.

To solve this problem, I appealed to the capacitarian account of control I developed elsewhere to argue that direct capacitarian control is retained by agents who unwittingly fail to fulfill a duty of care provided that they have the requisite abilities and opportunity to have complied with it and, consequently, they are directly blameworthy for this failure and indirectly so for the subsequent negligent action or

omission. I also explained how a commonsensical view of unexercised abilities allows this account to vindicate everyday attributions of culpability in cases of negligence and also showed in what respect my account is superior to Clarke's view.

Finally, I showed that the capacitarian account is able to rebut two important arguments that purport to establish that there is no culpability for negligence. Against Moore and Hurd (2011), I claimed that failure to exercise the relevant capacities can render an agent blameworthy for a negligent action regardless of whether such failure can be explained by a defect of the agent's practical rationality or character. And against King (2009), I claimed that the capacitarian account presented here offers the desired unified explanation concerning the blameworthiness of witting and unwitting wrongdoers. Thus, contrary to what these arguments purport to show, agents who inadvertently create a risk of harm, or inadvertently cause harm, can indeed be held morally responsible for their negligence.

References

- Alexander, Lawrence, Kimberly Ferzan, and Stephen Morse. 2009. *Crime and Culpability: A Theory of Criminal Law*, Cambridge: Cambridge University Press.
- Amaya, Santiago. 2022. "Negligence: Its Moral Significance", in J. Doris and M. Vargas (eds.), *Oxford Handbook of Moral Psychology*, Oxford: Oxford University Press.
- Brink, David and Dana Nelkin. 2013. "Fairness and the Architecture of Responsibility", *Oxford Studies in Agency and Responsibility*, Vol. 1. Oxford: Oxford University Press, pp. 284-314.

- Clarke, Randolph. 2014. *Omissions: Agency, Metaphysics, and Responsibility*. Oxford: Oxford University Press.
- _____. 2017a. "Ignorance, Revision, and Commonsense," in Philip Robichaud and Jan Willem
- Wieland (eds.), 2017, *Responsibility: the Epistemic Condition*, pp. 233–251. Oxford: Oxford University Press.
- _____. 2017b. "Blameworthiness and Unwitting Omissions," in Dana Nelkin and Samuel Rickless (eds.) *The Ethics and Law of Omissions*, pp. 63-83. Oxford: Oxford University Press.
- Hart, Herbert. 1968. *Punishment and Responsibility*. Oxford: Oxford University Press.
- Howarth, David. 2006. "Many Duties of Care: Or a Duty of Care? Notes from the Underground," *Oxford Journal of Legal Studies* 26: 449-472.
- Husak, Douglas. 2011. "Negligence, Belief, Blame, and Criminal Liability: The Special Case of Forgetting," *Criminal Law and Philosophy* 5: 199-218.
- King, Matt. 2009. "The Problem with Negligence," *Social Theory and Practice* 35: 577-595.
- Levy, Neil. 2011. *Hard Luck*. Oxford: Oxford University Press.
- McBride, Nicholas. 2004. "Duties of Care: Do They Really Exist?" *Oxford Journal of Legal Studies* 24: 417-41.
- McGeer, Victoria and Phillip Pettit. 2015. "The Hard Problem of Responsibility," in *Oxford Studies in Agency and Responsibility*, Vol. 3, pp. 160-188. Oxford: Oxford University Press.
- Moore, Michael S. and Heidi Hurd. 2011. "Punishing the Awkward, the Stupid, the Weak, and the Selfish: The

- Culpability of Negligence,” *Criminal Law and Philosophy* 5: 147-198.
- Murray, Samuel. 2017. “Responsibility and Vigilance”, *Philosophical Studies*, 174 (2): 507-27.
- Murray, Samuel, and Manuel Vargas. 2020. “Vigilance and Control”, *Philosophical Studies* 177: 825–843.
- Nelkin, Dana. 2016. “Difficulty and Degrees of Moral Praiseworthiness and Blameworthiness”. *Nous*, 50 (2): 356–378.
- Nelkin, Dana and Samuel Rickless. 2017. “Moral Responsibility for Unwitting Omissions: A New Tracing View”, in Dana Nelkin and Samuel Rickless (eds.) *The Ethics and Law of Omissions*, pp. 106-30. Oxford: Oxford University Press.
- Nottelmann, Nikolaj. 2007. *Blameworthy Belief: A Study in Epistemic Deontology*. Dordrecht: Springer.
- Raz, Joseph. 2011. “Responsibility and the Negligence Standard.” In Raz, *From Normativity to Responsibility*, pp. 255-269. Oxford: Oxford University Press.
- Rosen, Gideon. 2004. “Skepticism about Moral Responsibility.” *Philosophical Perspectives* 18: 295–313.
- Rudy-Hiller, Fernando. 2017. “A Capacitarian Account of Culpable Ignorance”, *Pacific Philosophical Quarterly* 98 (S1): 398–426.
- _____, Fernando. 2020. “Reasonable expectations, moral responsibility, and empirical data”, *Philosophical Studies* 177(10): 2945-2968.
- Scanlon, Thomas. 1998. *What We Owe to Each Other*. Cambridge MA: Harvard University Press.

- Smith, Angela. 2005. "Responsibility for Attitudes: Activity and Passivity in Mental Life," *Ethics* 115 (2): 236-271.
- Stark, Findlay. 2016. *Culpable Carelessness: Recklessness and Negligence in the Criminal Law*, Cambridge: Cambridge University Press.
- Sverdlik, Steven. 1993. "Pure Negligence." *American Philosophical Quarterly* 30: 137-49.
- Talbert, Matthew. 2017. "Omission and Attribution Error", in D. Nelkin and S. Rickless (eds.) *The Ethics and Law of Omissions*, Oxford: Oxford University Press, 17-35.
- Vargas, Manuel. 2013. *Building Better Beings: A Theory of Moral Responsibility*. New York: Oxford University Press.
- _____. 2020. "Negligence and Social Self-Governance," in Alfred Mele (ed.), *Surrounding Self-Control*. Oxford: Oxford University Press.
- Wolf, Susan. 1990. *Freedom within Reason*. Oxford: Oxford University Press.
- Zimmerman, Michael. 1986. "Negligence and Moral Responsibility." *Noûs* 20: 199-218.
- _____. 1997. "Moral Responsibility and Ignorance." *Ethics* 107: 410-26.

