

Negligence: its moral significance

Santiago Amaya
Universidad de los Andes

We often behave in reprehensible ways, sometimes knowing that we can do better. But, often we do it without deciding or choosing to do so, without even thinking that we risk doing something wrong. In other words, we behave *negligently*. Sometimes, specially if someone gets hurt, we get in trouble for it. Those who get hurt blame us; we feel guilty about it. Fortunately, if nothing catastrophic happens, we issue an apology and receive forgiveness in return. As far as social practices go, this seems to be a game with clear and easy-to-follow rules.

This chapter considers the *moral significance* of negligence. It does this by unpacking the complexity of the phenomenon and showing how it illuminates a range of related moral phenomena. As we shall see, negligence is not just a technical concept restricted to a narrow legal domain. Nor are attributions of it a marginal social and moral phenomenon. Negligence, in fact, is a test case for questioning assumptions central to widespread ways of theorizing about moral responsibility. It is also a window into everyday practices behind judgments of culpability, attributions of blame, and allocations of punishment.

I begin with a statement of what negligence is, a specification that illuminates why it is morally significant. As it will become evident, there are many ways in which negligence can occur. So, after defining negligence, I explore this variation, explaining why it is unlikely that a single approach can handle all cases and detailing some problems with less over-arching existing approaches along the

way. In the end, I discuss some reasons for being skeptical of the moral significance of negligence and conclude with some reflections on the prospects for future work.

1. Voluntarism

Traditionally, the domain of moral responsibility has been seen—some theorists think it ought to be seen—as circumscribed by the *voluntary*.¹ There are various ways of understanding what makes a piece of behavior voluntary, corresponding to how one thinks about the motivational and cognitive aspects of voluntary conduct. So, the view has been formulated in a variety of ways. A common formulation, however, emphasizes the presence of intentions as a central criterion for the attribution of responsibility. According to it, people are responsible only for their *intentional actions* (or omissions) and the *intended or foreseen consequences* of them (Zimmerman 1988; Levy 2014).

Negligence occupies a middle ground between *intentional* and *merely accidental* wrongdoing. Yet, it seems reasonable to hold people accountable and to blame them for their negligent conduct. That's why many theorists working on

¹ The view that the voluntary defines the domain of responsibility can be traced to Aristotle's discussion in book III of the *Nicomachean Ethics*. But it is also held by contemporary theorists in a variety of forms. Besides the accounts couched in terms of intentional actions, there are accounts of it in terms choice (Duff 1993; Moore 1997: 548-592) and, as I argue below, also in terms rational control (Wolf 1990; Fischer and Ravizza 1998)—although the latter are often thought to be an alternative to voluntarist approaches.

moral responsibility and blame have recently turned their attention to it.² Being an example of wrongdoing that is not intentional or foreseen, but cannot be discounted as merely accidental, negligence is a *prima facie* challenge to the *voluntarist* view.

Negligence has figured prominently in discussions within Anglo-American tort law, at least going back to the mid 19th century.³ But this should not lead us into thinking that negligence is just a technical concept in the domain of institutionally sanctioned reparation. Negligence, in fact, is part of the machinery by which many adults carve their social world, attribute responsibility to each other, and think about the allocation of punishment (Shultz and Wright 1985; Nuñez et al 2014). Although voluntary wrongdoing might be the stereotypical target of attributions

² Among philosophers, much of the recent enthusiasm for negligence derives from George Sher's (2006) provocative paper "Out of Control." Since then, there has been a large amount of new work on negligence, way more than can be fairly referenced here. Two recent volumes dedicated to the topic of culpable ignorance are worth mentioning, Robinchaud and Wieland (2017) and Nelkin and Rickless (2017a), as they present up-to-date versions of some positions developed over the last decade. An excellent survey of the philosophical intricacies of some of this discussion can be found in Rudy-Hiller (2018). In 2011, *Criminal Law and Philosophy* ran an issue dedicated to the legal aspects of negligence. For references, on the psychological investigation of judgment of negligence see below.

³ The case that is said to introduce negligence as a distinctive tort is 1850's *Brown vs. Kendall*, 60 (Mass) 292, where the reasonable man's test makes one of its first appearance in US legal history.

of responsibility, our everyday blaming practices clearly extend beyond the stereotype.

This is something worth emphasizing. The traditional wisdom among developmental psychologists going back to Piaget (1932) and Kohlberg (1969) is that the development of moral judgment is characterized by an increased *shift from outcome to intent* (for contemporary statements of the view, see Helwig et al 2001, Zelazo et al 1996). Recent studies, however, have provided evidence showing that this might not be as radical a shift as it was previously hypothesized (Hamlyn et al 2007; Cushman et al 2013;). In particular, when stimuli are carefully controlled for intention salience and recency, young children seem to be more sensitive to intent than what they otherwise would seem to be (Nobes et al 2016, 2017).

Negligence provides a complement to this story, but from the opposite end of the developmental spectrum. It is not just, as some psychologists have emphasized (Young et al 2007; Cushman 2008; Alicke 2014), that automatic evaluation by adults tends to be more outcome-focused, whereas controlled evaluation is more intent-driven. In addition, and perhaps even more interesting, what the study of negligence indicates is that even the *considered and deliberate* moral judgment of adults is not as reliant on intent as the traditional picture presupposes. In fact, when it comes to negligence, the moral judgment of older adults seems to exhibit a shift *from intent to outcome* (Margoni et al 2019).

It might seem, at least when considered in the abstract, that these two issues—the success of our standing theories and the form of some current practices—are independent of each other. Even if culpable negligence in the end proves voluntarism to be false, it is still an open question whether, say, negligent agents in the real world are blamed for the right reasons or should be punished in

accordance with lay attitudes. Still, unless these reasons and attitudes are independently found problematic, their prevalence does serve to underscore the challenge outlined here. At least, it suggests that the challenge raised to voluntarism goes beyond the zero-sum game of building theories and offering counter-examples to it. Understanding to what extent and why it is reasonable to hold negligent people accountable is part of understanding *the moral legitimacy* of some of our current social practices.

2. Fixing the reference

Negligence, as said above, differs from intentional and merely accidental wrongdoing. To a first approximation, the contrast goes something like this. A person can behave negligently by intentionally doing many things. But, in so far as her behavior is truly negligent, the wrongdoing should not be viewed as intentional or as foreseen. At least, it should not be viewed as intentional or foreseen *under the description* in which it is also wrong. The negligent agent, in other words, does not act intending to do wrong, nor does she intentionally or knowingly do wrong for the sake of anything else.

A person's behavior, on the other hand, can fall below some acceptable moral standard due to its unintended and unforeseen consequences. The consequences might have been easily preventable, even by the person herself. All of this, however, is *compatible* with the episode being a mere accident (say, a product of bad luck) and the person not being culpable for it. Negligence, in other words, does not just imply getting things wrong, where there was the option of getting them right. It requires, in addition, the presence of a *mistake*.

To sharpen the contrast, we need to distinguish negligence from other forms of wrongdoing that are superficially similar but that do not constitute a challenge to voluntarism. Consider to this end the following hypothetical scenario:

Cooking. Annie invites a friend to come over for lunch on Sunday. He accepts but warns her about his peanut allergy. Annie figures making her pesto recipe would probably be a bad idea, so she decides instead to make chicken. On Sunday, running behind schedule, Annie fries the chicken using some left over peanut oil, which she finds in her kitchen. Sure enough, shortly after the friend starts eating he has an allergic reaction to it.

(Adapted from Nuñez et al 2014, study 3)

Many people would find Annie responsible for the allergic reaction of her friend and would blame her for it (Nuñez et al 2014). But there are different ways of understanding the story. And not all of them are instances of negligence, at least in the sense in which negligence is meant to create trouble for voluntarism.

It is possible, for example, to construct the scenario as one in which Annie makes a *reckless decision*, as opposed to a negligent one. On Sunday, she remembers her friend's allergy, she knows the oil contains peanuts, but is desperate to have lunch served in time. So, even though she does not intend to harm her friend, she goes ahead and cooks with the peanut oil she has in hand. Annie simply hopes her friend won't be harmed by it.

The story can also be expanded in ways where the harms occur due the agent's *excusable ignorance*. Suppose, for instance, that Annie used the oil without knowing that it was peanut oil. Unbeknownst to her, her roommate had by mistake refilled the olive oil container with peanut oil a few days before. Then, even though the harm could have easily been avoided (Annie could have not used

the oil, she could have checked before using it, etc.), it does not seem that Annie acted out of negligence.

It should be clear why these variations of the story are *compatible* with the truth of voluntarism. Reckless Annie might not do wrong on purpose, but she knowingly *risks* doing it. Ignorant Annie, on the other hand, might have informed herself better. But, in so far as there is no reason for her to suspect that she was misinformed, her lack of knowledge seems *not* to result from *her own fault*. As we shall see, some usages of the term “negligence” count reckless decisions as instances of it. But, for the discussion in this chapter, we shall not be concerned with these sorts of cases. We shall restrict ourselves to instances of what might be called, perhaps redundantly, *inadvertent negligence*.

Even this restricted sense, negligence can occur in multiple forms. And, as we shall see, it is important to keep this variability in mind when trying to explain why negligent actors are held culpable. Annie, for instance, might have *negligently decided* to use peanut oil, say, if she didn’t remember her friend’s allergy when she started cooking on Sunday. Alternatively, she might have *negligently failed to realize* that being allergic to peanuts meant being allergic to peanut oil. The issue of the allergy did occur to her but she simply didn’t connect the dots. Finally, Annie might have *negligently omitted* checking whether the oil contained peanuts, which she should had, given that her roommate had several times in in the past refilled the container with peanut oil.

These variations of the story differ from each other in important ways. For instance, in the first and second version of it, the negligent agent makes a bad decision. The last one, by contrast, does not involve making a decision at all. It is, instead, a pure omission. Likewise, whereas the second version of the story

involves a mistake regarding the *truth* of Annie's beliefs, the mistake in the first and third versions seems to be located elsewhere: respectively, a failure to remember certain things already believed and a failure to seek relevant information.

For the moment, however, we can put these differences aside and focus on the superficial shape shared by all of them. The negligent person has enough information to act in line with certain moral standards or, at least, has easy access to information that would allow her to do so. But because she *fails to advert* to that information or to the need to seek it, she winds up acting below those standards. What she does, therefore, is not a reflection of what she intends or how willing she is to do or to risk doing things with potentially bad consequences. It is a rather a reflection of her failure to see her actions (or her omissions) in the light of some available information at the time.

3. Standard definitions

Legal scholars commonly define negligence in terms of a failure of *due care or diligence*. This characterization follows the definition of negligence established in the Moral Penal Code (Owen 2007; Raz 2010, Shiffrin 2017). The animating thought is that people have a general obligation to behave with enough care, so as not to harm others or risk harming them without any reasonable justification. Negligence is constituted by a breach of this obligation. The negligent agent harms or risks harming someone because she fails to exercise the care expected from her. Her inadvertence is supposed to be the psychological manifestation of this failure.

It is easy to see how this general duty to care works when institutionalized relationships are involved. Parents, for instance, are under the duty to look after

the wellbeing of their offspring and to make *well-informed* decisions when it comes to them. It is no surprise, therefore, that illustrations of negligence in the law often involve cases of poor parental decisions. It is no surprise either that real-life cases of parental negligence more often than not result in convictions (Collins 2006).

It might be harder to see how this kind of general duty works when there is no special relation between the agent and her potential victims. But it is still plausible to think that this *deontic framework* applies here too (see McBride 2004, Howarth 2006, Owen 2007 for discussion). Even if there is no codified duty to care for pedestrians or other drivers in the road, one is under the obligation to drive carefully. In general, the absence of a specific duty to care for others is compatible with the existence of a duty to *avoid carelessness* when performing certain risky activities.

Now, as mentioned earlier, the *concept* of negligence is not restricted to civil or criminal courts—or to philosophical discussions concerning moral responsibility, for that matter. Not only do people in everyday life seem to respond negatively to negligent behaviors and to treat them as legitimate targets for attributions of culpability and grounds for punishment. But they also seem to have some sort of *explicit mastery* of the concept. This has, in fact, led some scholars to insist that the legal category of negligence is borrowed from popular morality, although it is unclear how much evidential support there is for the view.⁴

⁴ A statement of this position, which is intended as common knowledge among legal theorists, can be found in Raz (2010: 5). The view that many legal distinctions and principles are just codifications of common-sense morality has a long tradition

At any rate, empirical work concerning everyday conceptions of negligence does suggest that some of the deontic structure found in the law can be also found within the folk conception of negligence. Participants in these studies often contrast negligent with intended and desired behavior (Shultz and Wright 1985; Nuñez et al 2014, Laurent 2016, study 3). When freely asked to define negligence, a prevalent response is that negligent behavior is careless or a sign that the agent fails to take reasonable care (Nuñez et al 2014, study 2). Attributions of negligence are sensitive to the lack of precautions surrounding the act (Karlovac and Darley 1988).

Most of these studies, unfortunately, have presupposed too broad an understanding of what negligence is. So, some of its conclusions need to be taken with a grain of salt. Some of the landmark studies, for instance, ask participants to rate the culpability of an actor for “negligent” wrongdoing without explicitly defining the term. But the scenarios evaluated lack enough detail to rule out the possibility of the actor being recklessness, as opposed to negligent (Shultz and Wright 1985; Nuñez et al 2014, study 3). Others studies provide participants with definitions of “negligence” and ask them to rate an actor along that dimension. But sometimes the definitions provided are compatible with conscious wrongdoing (Nuñez et al 2014, study 1; Laurent et al 2016, study 1) or, to make matters worse, they explicitly include the possibility of deliberate wrongdoing (Karlovac and Darley 1988).

among psychologists. See, for instance, Hamilton (1978) and Fincham and Jaspers (1980).

Of course, it is possible that lay people are not sensitive to these distinctions. Or perhaps, they do not mark the difference with the terminology that is standard among legal scholars and philosophers.⁵ But this is *not* something that the existing studies demonstrate, as opposed to something that they presuppose. The truth is that at present, we do not have enough evidence to determine what are the exact boundaries, if any, of the lay concept of negligence.⁶

⁵ It is possible, of course, that the distinction between negligence and recklessness is too fuzzy to serve certain kinds of purposes. For arguments that seek to demonstrate that the distinction is legally unstable, see Simmons (2009) and Husak (2011).

⁶ It does not help in this regard that plenty of the available data about conceptions and attitudes towards negligence is buried under studies that purport to be about other things. Lagnado and Channon (2008), for instance, study unintended foreseeability evaluations with respect to unintended wrongdoing. Their materials are mixbag of negligence stories vs. stores of mere accidental wrongdoing. Cushman (2008) discusses some of his results, as concerning the moral evaluation with regard to accidents, but some of the vignettes he uses clearly involve negligent conduct. Young et al (2010) present their study as one about moral luck, when in reality it concerns the effects of belief evaluation in some forms of negligent behavior. This is meant not a criticism of their results, but as an indication of how much data about folk conceptions of negligence there could be that we haven't properly thought about because it hasn't been properly filed and categorized.

4. A guilty mind?

Voluntarism restricts the domain of morally responsibility to voluntary wrongdoing. Above we defined the voluntary in terms of our intentional actions (or omissions) and their foreseen consequences. But that is obviously only one way, among others, of characterizing it. Alternatively, the voluntary could be defined by appeal to choice, rational control, or other similar notions, as long as these provide motivational and cognitive analogs of intent.

There are, as Annie's story of exposing her friend to peanut oil illustrates, various ways in which people can behave negligently: they can *forget* important information, *miscalculate* risks, or even *fail to see* the danger in conducts that were evidently risky. To the extent that these mistakes result in distinct forms of wrongdoing, the relevant lessons drawn from them apply *mutatis mutandi* to other ways of defining the voluntary. Episodes of negligence, once one starts thinking about them, also problematize accounts of responsibility in terms of choice, rational control, etc.

Let me explain where I think where the problem comes from. Voluntarists of all stripes conceive of culpability in terms of what might be labeled a *guilty mind*. The culpable agent, on their view, need not act out of malice, that is, motivated to do wrong as such. Her intentions, decisions or choices need not be, in and of itself, morally censurable. Yet, at a minimum, she must have been motivated to do what she did in its wrongful guise. That is, she must have regarded the action, maybe

not as wrong, but at least as having the characteristics that ultimately make it wrong.⁷

In cases of negligence, however, this wrongful characterization is absent from the agent's mind. This is what the failure of advertence comes to. The negligent person seems to be at fault for the wrong she did. But at the time she does not see her behavior as having the features that made it the instance of wrongdoing that it turned to be. Hence, *if she is indeed culpable, her culpability cannot be grounded on her having a guilty mind at the time.*

Now, as intimated above, the problem for voluntarism is that widespread attitudes towards negligence suggest that the antecedent of this conditional is true. In particular, negligent agents are not just regularly considered distracted, ignorant, careless, etc. That is, these negative attitudes cannot be understood merely as aretaic evaluations. Nor can they be interpreted just as expressions of frustration, or disappointment, although negligent wrongdoing is, no doubt, at

⁷ A succinct way of putting this perspective on culpability is offered by Levy (2014: 37), who argues that blameworthiness does not require a *de dicto* belief on the wrongness of the actions (e.g. what I am doing is wrong) but a *de re* belief about the action's wrong making features (e.g. I am using an ingredient to which my friend is allergic). See also Duff (1993) and Moore and Hurd (2011: 150). Another way of casting the view of culpability as a guilty mind is in terms of the person's having a deficient quality of will. The way traces back to Kant, who thought the will of a person was both the locus of her moral worth and her moral *appraisal*. For a contemporary formulation of the view of culpability as bad quality of will view, see Mckenna (2011).

times frustrating and disappointing. At least in some cases of negligence, agents are actually judged to be culpable for what they did.

Legal precedent, in so far as it reflects our practices of moral evaluation, is telling in this regard.⁸ And so are our everyday spontaneous responses to episodes of negligence. Consider one last time the case of Annie. If cooking with peanut oil was truly negligent (nor intended or reckless) of her, then perhaps it might be too much for her friend to resent her: she didn't trigger his allergy on purpose or risk hurting him to have lunch ready on time. Still it would not be unreasonable for him to think that the allergic reaction was her fault and even to expect her to feel bad and to apologize for it. The same would be true about Annie herself. She would probably regret the whole thing, feel guilty about it, and be willing to do some things to earn forgiveness. That is, both Annie and her friend would have had the kind of reactions that are best explained in terms of a judgment of culpability.

This is not just something that anecdotal evidence suggests. The scientific probing of attitudes in studies where participants are presented with negligence scenarios also provides evidence for it. Typically, in these studies actors are found morally responsible for their wrongdoing, less responsible than for voluntary offenses but responsible nevertheless (Shultz and Wright 1985). They are considered blameworthy or guilty for the harms inflicted (Kneer and Machery

⁸ Some legal theorists have argued that negligence should disappear as a category of culpability (Moore and Hurd 2011, Alexander and Kessler Ferzan 2009, Finkelstein 2005). But these theorists are obviously endorsing a revisionary position.

2018; Murray et al 2018); in fact, they are considered blameworthy also when no harms ensue (Young et al 2010). Lastly, for offenses of varying degrees of seriousness actors are found deserving of punishment, even obligated to pay restitution (Karlovac and Darley 1988, Laurent et al 2016),

Judging by these reactions, it is clear that voluntarism does not offer a *realistic description* of our existing practices of attribution of culpability. People are indeed held responsible for their voluntary behaviors, but not *only* for them. This lack of alignment between theory and practice can be interpreted in different ways. It might give us a reason to revise the practice in the light of the theory, that is, to regiment our common attitudes towards negligence in accordance to it. Alternatively, it could be grounds for making amendments in the theory, superficial or deep, depending on how big one thinks the negligence challenge is.

5. Revision

Most theorists working on negligence have argued that these practices *ought* to be vindicated. So, they agree that some degree of theoretical revision is appropriate in this respect. Voluntarism, as far as they are concerned, does not just fail to provide an accurate description of some current practices but it is normatively inaccurate as well. There are, according to this view, *moral reasons* for expanding the domain of responsibility beyond the voluntary.

Two types of considerations have been prevalent in this regard. There are, first, considerations of *fairness* (Amaya and Doris 2015; Clarke 2014; ch. 7; Raz 2010). In brief, in so far as agents have a fair opportunity to behave well, it would seem appropriate to judge them culpable for not behaving as they should. The mistake, to be sure, might mitigate the amount of blame they deserve. But it does

not provide an excuse for the fact that they didn't acted as they we supposed to, when it was not onerous for them to do so.

There are, on the other hand, considerations having to do with our *moral cultivation* (Pereboom 2015; Vargas forthcoming). If part of the rationale of holding each other responsible for certain forms of wrongdoing is that this encourages and helps sustain certain forms of moral agency, then it would seem reasonable to include negligent wrongdoers within the scope of these practices. Blaming the negligent, in particular, not only serves as a warning sign. But it also expresses a collective commitment to treat certain forms of carelessness as problematic for social life.

Both of these considerations clearly admit of a sharper formulation. To take the point about cultivation as an illustration, the negligent actor is by definition not aware that she is being careless at the time. So, the general injunction to be careful codified in our adverse attitude toward negligence will be an unlikely guide at the moment in which the negligent act can occur (Alexander and Kessler Ferrzan 2009: p. 273-274). If it ultimately shapes her behavior, it would probably by introducing changes in her general routines and habits so as to avoid possible episodes of inadvertent wrong doing.

Obviously, one thing is to argue that negligence is morally appropriate grounds of culpability. It is another to say *what grounds* the culpability of the agent in these cases. The latter, as we shall now see, is a much harder enterprise. It is not easy to see what could replace the account of culpability in terms of a guilty mind, which ultimately underwrites voluntarism. Culpability, after all, is something that attaches to a moral agent, not to an act. And a moral agent, it would seem, is hardly dissociable from the moral contents of her mind.

6. Variability

Negligence is a thin notion. It is a determinable that obtains in many different determinate ways (Tappolet 2004). This is, in fact, one of the reasons why it is so hard to explain what grounds our culpability for negligence. Proposals that do a relatively good job explaining why some negligent agents are appropriately held culpable often leave unexplained significant swaths of negligent behavior. Or, they manage to account for some episodes of negligence at the cost of making the phenomenon less morally pervasive.

There is, however, as I now turn to discuss, some systematicity within this variability. Although negligent episodes can occur in multiple ways, there are some well-defined dimensions along which the variation occurs. Looking at these axes of variations shows how complex the psychology behind negligent behavior can be.

6.1. Competence vs. performance

One dimension of variation, already mentioned above, concerns the type of inadvertence exhibited by the agent. Here, at least two broad possibilities can be distinguished. Some negligent agents inadvertently do wrong because at the time they *fail to know* or to believe certain things that they should. Others know and believe what they should but they *fail to bring it to mind*, so they inadvertently wind up behaving contrary to it. In other words, whereas some negligent agents should have known better, others actually know better.

In so far as both are forms of negligence, each involves some sort of *epistemic shortcoming*. The person, after all, acts failing to see that her actions (or omissions)

are wrong. Still, there is an important *psychological* difference among them. Some negligent mistakes impugn the agent's moral competence. Obviously, they do not put into question the person's ability to perform the type of action that would allow her to discharge her obligations. But they put into question her capacity to appreciate that a certain moral standard applies in general, or in the particular circumstances she occupies. At least, her not knowing something that she should have known is indication of her being morally incompetent in some respect.⁹

Other mistakes, by contrast, are simple failures to exercise that competence. They are, in other words, pure *performance mistakes*. The person, at the time, does not have a moral impairment preventing her to see what the right course of action is. Instead, she fails to see that she is doing wrong because some of her knowledge or beliefs fail to become active in her mind. In consequence, they do not alert her that a certain course of action needs to be taken or that a course of action already chosen is somehow morally inappropriate.¹⁰

⁹ Marcia Baron (2001) discusses cases of this sort, involving rape in the light of a false belief about the other part consenting to sex.

¹⁰ The distinction between errors due to a lack of competence and pure performance errors is familiar from psycholinguistics, since Chomsky (1964; 1965) introduced it to criticize behaviorist approaches to linguistic cognition. But clearly the difference is not restricted to our linguistic performances; it is applicable to any domain where some form of knowledge (tacit or explicit) needs to be implemented under conditions of limitation and time pressure. For discussion of how the distinction works in the moral domain, see Amaya and Doris (2015) for discussion.

The difference is, perhaps, best brought to light by considering *what would remedy* in each case the person's epistemic shortcoming. In one type of case, remedying it would require providing her with information she does not have at the time, say, information about the moral obligations that apply to her, the unforeseen consequences of her actions, or the moral significance of them in the context of action she occupies. In the long run, this might require developing *moral sensibilities* that she lacked at the time. In the other type case, by contrast, it would simply be a matter of pointing at information that she already has or developing *cognitive routines* for making certain possessed information more available. To behave well, in other words, one would not have to modify the *contents* of the agent's mind but make more salient what she already knows.

This psychological distinction marks an important *moral difference*. It is one thing to hold someone accountable for violating a moral standard, while her competence to meet that standard is admittedly compromised. That would seem to run against the considerations of fairness mentioned earlier.¹¹ But it is an altogether different thing if her negligence is purely due to a performance mistake. For there, by definition, the person has intact the competence to behave well, so she can serve as the standard for her own evaluation. Given what she knew or believed at the time, she should have definitely acted otherwise.

¹¹ Zimmerman (1997) and Rosen (2004) both make the point in terms of false beliefs: it is unreasonable to expect that a person will behave contrary to what she believes. But the point can also be made about agents who simply lack the relevant beliefs: how can someone be expected to act as she should, if she does not know what she should do?

Negligence theorists focusing on performance mistakes have accordingly favored “capacitarian” approaches (Clarke 2014; Murray 2017; Rudy-Hiller 2017). On their view, culpability need not be grounded on voluntariness or, more broadly, on the agent having a guilty mind. It could alternatively be grounded on the person having the capacity to behave as she should, where the capacity is a matter of her intrinsic abilities and the environment cooperating with her.¹² In essence, to the extent that the negligent agent is competent enough at the time of her mistake, it is fair to hold her responsible for it. It does not matter that she was not aware at the time of the wrongness of what she was doing.

In so far certain forms ignorance undermine a person’s moral competence, it is clear this kind of appeal to capacity will not work for them, at least not without some serious modification. Instead, for these cases, theorists have often invoked *tracing strategies*.¹³ The general idea here is that culpability in this sort of case is

¹² The capacitarian approach can be read in two different ways, a weak and a strong way, although these two readings are not always clearly distinguished. One could think that capacity is, alongside with others, a sufficient condition for moral responsibility (see Clarke 2014: 167). Alternatively, one could think that capacity or fair opportunity is what grounds moral responsibility in general, whether this applies to voluntary or involuntary wrongdoing (see, for instance, Brink and Nelkin 2013).

¹³ Seminal discussions of tracing strategies can be found in Fischer and Ravizza (1994) and Zimmerman (1986, 1997). But see also Rosen (2004), Levy (2009) and, to some extent, Nelkin and Rickless (2017b). Here I discuss voluntarist-friendly varieties of tracing, but there are non-voluntarist versions of it in Fitzpatrick (2008),

derivative, as opposed to basic. Perhaps, at the time the agent was not competent enough to act as she should have. Yet, if her incompetence traces back to a moment in time where she had enough information to anticipate the mistake, things arguably look quite differently. The culpability for the wrongdoing could *depend* in that case on the culpability of the prior act.

Traditionally, tracing strategies have followed the spirit, if not the letter, of voluntarism. For them, the culpability of ignorant wrongdoing can be traced to a *benighting* intentional act: a prior moment in which an agent intentionally or knowingly let pass the opportunity of improving her cognitive position or in which she made a decision that actually impaired it (The term “benighting” comes from Smith 1983). Tracing theorists, in other words, substitute the requirement of the agent’s mind being guilty at the time with the requirement of her ignorance resulting from the agent’s prior guilty mind.

It is doubtful, however, that this strategy works *as an account of negligence*. To begin, it is unclear to what extent paradigmatic instances of negligence are in fact preceded by a benighting act, as the theory requires (for discussion see Vargas 2005, Smith 2011). As critics have pointed out, in some of these cases, the agent’s actual ignorance is simply not traceable to an earlier decision made by her. When it is traceable, the earlier decision often does not seem culpable (But see Fischer and Tognazzini 2009 and Timpe 2011 for a rejoinder). And, even if it is,

Shabo (2015). In particular, Murray (2017) and Rudy Hiller (2017) appeal to tracing in ways that are decidedly contrary to voluntarism. Tracing is, obviously, not the only way of accounting for this sort culpable ignorance. Talbert (2013) and Harman (2017) offer attributionist accounts of the sort that I discuss below.

phenomenologically speaking at least, our judgments of culpability in these cases do not seem to be a direct response to it. We arrive to our judgments of culpability, or so it would seem, without (or prior to) finding a tracing anchor (Graham 2017).

More significant, perhaps, it is unclear whether tracing actually gives an account of culpable negligence, as opposed to *merely dispensing with it*. For it seems that what these proposals ultimately do is to deal with certain cases of negligence by reducing them to cases of *direct responsibility* for recklessness (Agule 2016; King 2017). In the resulting account, the agent's culpability is just the culpability of having knowingly taken some unjustified risks, in particular the risk of landing in a suboptimal epistemic position. Given that, holding the agent culpable for what she did afterwards seems moot. Her acting out of ignorance later seems not to add anything.

6.2. Moral vs. factual

A second dimension of variation concerns the type of information to which the agent fails to advert and *on account of which* she is said to be negligent. Here, there would seem to be two major dividing lines too. The agent, to put it briefly, could fail to see that some moral obligation applies to her. Alternatively, she could fail to see some relevant aspects of what she is doing, the circumstances she occupies, or the consequences of her actions given those circumstances.

The distinction is sometimes drawn in terms cases that involve neglect of some *moral principle* vs. those involving neglect of *factual information*.¹⁴ But drawing the contrast in these terms is a bit tendentious. This is not just because moral

¹⁴ See, for instance, Rosen (2004), Guerrero (2017), Wieland (2017: 2).

principles, one could argue, are grounded on *factual claims*—at least, many ethical realists have argued that this is so. But because it runs the risk of obfuscating one interesting way in which many negligence cases can be plausibly analyzed.

To see this in more detail, consider the following case:

Bathing. John is at home giving his 2-year old daughter a bath. He fills the bath, while her daughter stands near the tub. The phone rings in the next room. John tells her daughter to stand near the tub while he answers the phone. John believes her daughter will stand near the tub for a few minutes and wait for him to return. When he returns, his daughter is in the tub, dead, face down in the water.

(Slightly modified from Kneer and Machery 2018)

Many people would find John culpable for what he did (Kneer and Machery 2018).¹⁵ It was, we can agree, a matter of bad luck that the child got into the tub instead of waiting for her dad. But deciding to leave the child alone next to it was a negligent mistake. Admittedly, John didn't think that doing it would put his daughter under the risk of drowning. But definitely he *should have thought* about it.

It would seem, thus described, that John's inadvertence concerns a factual matter: the likelihood of his decision having a certain outcome. But this is not the only way of describing what John should have considered then. There is at least an

¹⁵ Different variations of the case have been discussed among negligence theorists (see, for instance, Alexander and Ferzan 2009; Husak 2011). For present purposes, we can stipulate that the case involves a negligent (not a reckless) decision, which is not necessarily due to a memory failure. The parent's mistake occurs at the time she decides to leave the child next to the tub.

alternative description of the case where John's neglect does not concern the factual aspects of his situation, a description that still falls short of making a moral principle the object of his neglect. To wit, John *should have realized* that children *ought* not to be left alone next to a bathing tub.

Superficially, the two readings, the factual and the one just proposed, come down to the same thing.¹⁶ However the case is described, it is true that John did not anticipate that the child could drown; given the outcome, probably he should have. But as an account of what makes him negligent there is a significant difference between the two. In one reading, John is made negligent by his failure to *foresee* the consequences of his decision. In the other, what makes him negligent is his failure to appreciate a *well-known rule* for taking good care of a child.

The rule, it is worth emphasizing, is not a moral principle. There is nothing inherently wrong with leaving a child unattended next to a tub filled with water, as there is, say, in breaching the obligation not to kill or not to lie. In fact, there is nothing wrong at all with breaching the rule, as long as enough precautions are taken to avoid bad consequences from ensuing. Like other similar rules, this is just

¹⁶ Legal scholars often distinguish between two customary tests of negligence (For discussion, see Hurd and Moore (2002) and Owen 2007)). First, there is the well-known foreseeability test: *was the harm foreseeable to the defendant when she acted?* There is also the lesser-known harm-within-risk test: *was the harm of the type that motivated the prohibition of the defendant's conduct?* This second test is akin to the analysis proposed here, although to be strict the HWR test is often viewed as a test for proximate causation, which is only one of the components of the negligence tort.

a maxim useful for codifying some of the practical obligations that follow from a general moral principle—in this case, the principle is that parents ought to care for their children. In and of itself, there is nothing “moral” about its content.

Among negligence theorists, at least when it comes to cases of this sort, the *foreseeability* approach has been almost ubiquitous. Actually, because the likelihood of the bad consequences of a decision might be offset by its positive consequences, this approach ultimately views the culpability of the agent as residing in her departure from a standard of risk calculation (Feinberg 1987: 190ff; Timpe 2011). Here a common analysis is the reasonable person test, according to which the standard of risk calculation is fixed in relation to what a reasonable person would do in an analogue situation.

As common as it is, this approach is hostage to a series of interrelated *normative* objections. It is not clear who the reasonable person is supposed to be. Even if it were, it is not clear why moral agents should be held to this standard. Intuitively, being poor at making risk calculations might be less than ideal; it might be a sign that one has some cognitive shortcoming. But it does not seem to be something that makes one deserving of moral blame (for detailed criticisms of this test, see Zimmerman 1986, Hurd and Moore 2002; Alexander and Ferran 2009; Sher 2009: ch 5).

There are, in addition, *psychological* reasons to doubt applications of the test, as a way of determining whether someone else is culpable for not advertent to the risk she was creating. Even if the normative objections can be handled, it seems that people are consistently poor in estimating the foreseeability of a certain outcome obtaining. We are, in short, prone to exhibit a hindsight bias: knowing

that the outcome obtained has a clear effect in our estimations of the probability that it could obtain.¹⁷

Considered in the abstract, this might not be an unreasonable heuristic; absent better information, the fact that an event happens might be a good reason to judge it more likely than what one would otherwise think. The problem is that what matters for negligence attributions is *ex ante*, not *ex post*, foreseeability. That is, it matters whether at the time in which the agent acted and with the information she had, the likelihood of the bad consequences were low enough. Given this restriction, using information about the outcome obtaining to make this assessment is certainly a mistake.

6.3. Lack of care vs. carelessness

Negligence involves a failure of due care. Failing to care, as we have seen, excludes the intent or the active desire to breach some moral obligation. But that still leaves unresolved an ambiguity in the definition that needs to be cleared up.

Corresponding to it, there is a third dimension of variation in negligent behavior worth taking into account.

To see it, consider the following scenario:

Groceries. Randy's wife calls to ask if he can buy some groceries on the way back from work. She needs them for a party that she is hosting in the

¹⁷ The original finding is reported in Fischhoff (1975). Roese and Vohs (2012) provide a recent review of the findings. For a documentation of the hindsight bias as affecting judgments of negligence, see Kamin and Rachinsky (1995) and LaBine and Labine (1996).

evening. Later that day, Randy takes the usual route planning to make the stop at the store. On the way, however, he gets engrossed on his own thoughts and winds up driving straight home, inadvertently failing to make the stop. By the time he sees his wife and realizes it, the store is closed.

(Modified from Murray et al 2018)

Randy, we can stipulate, does not want to secretly spoil the dinner plans or to upset his wife. He actually makes a plan to stop at the store and get the things his she asked for. But as the time approaches, his mind begins to be filled with thoughts about work. As a consequence, he inadvertently omits the stop and winds up driving straight home.

Still, there various ways of understanding the omission, compatible with all the “facts” mentioned so far. First, perhaps Randy forgot to stop at the store because he did not *care enough* about the dinner his wife was hosting that evening. It is not necessary to assume here that he felt some sort of animosity toward the event or, for that matter, toward his wife. It is just that, as far as his priorities go, getting everything ready for the evening, as he had committed to do, was not at the top of his list. That’s why he didn’t even think about it as he drove past the store.

Another possibility is that Randy simply failed to *exhibit how much* he cared about it. Let’s suppose that he cared enough about the dinner, or at least that he cared as much as anyone normally cares about these things. Yet as soon as he got into his car and started driving on the familiar route, he automatically defaulted to the routine of driving straight back home. Whatever his attitudes were towards his wife’s dinner plans for that evening, they didn’t shape what he did.

These ways of telling the story illustrate two ways in which a person's behavior can be assessed as *careless* and, therefore, as negligent. Sometimes, what the person does is made careless by the *motivational structure* that brings it about. Her behavior, in those cases, reveals how the relevant attitudes that motivate her to act, comparatively speaking, fare with respect to each other. They reveal, for instance, how little the person cares about dinner plans, or about pleasing his wife, in comparison to what he cares about other things.

Other times, however, what the agent does is careless simply because her conduct fails to align with certain maxims of care of the sort that were mentioned above (e.g. do not leave children unattended, if you offer to do things in the future, set reminders that would bring those commitments to mind, etc). In such cases, there might some attitude of the agent revealed by her actions or her omissions, say, a desire to get home and relax. Yet what she does is *not a reflection of the balance* of her motivations of the time. The agent had an intention to do something. All things considered, there was nothing she preferred to do then. She didn't change her mind about it. But the intention slipped her mind and she wound up doing something else.

Slips of this sort are actually common in everyday life (Norman 1981; Reason 1984, Amaya 2013). It is only because their consequences tend not to be too bad that we underestimate the frequency of their occurrence. But it is enough to think about a routine and/or boring activity in one's life to bring episodes to mind that remind us of how often they happen. Perhaps, it is not the groceries that one fails to pick up. It is the library book that you fail to drop off, the exit at the freeway that you drive past, the glaring typos in your manuscript that you let stand.

Various factors potentially explain why slips like this happen. Described in terms of what the agent failed to do, they look like memory lapses: he forgot to pick up the groceries. Described in terms of what he did, they look instead like attentional captures: he was too distracted by his thoughts about work. One interesting hypothesis is that these slips are *errors of vigilance* (Amaya and Murray 2017). They result from the erroneous allocation of memory and attention necessary for the pursuit multiple interrelated goals (say, to think about work, while driving to the store, on one's way back home) in circumstances where the likelihood of a mistake is perceived as low.

Again, this psychological difference is *morally significant*. Various philosophers influenced by Kant have put forward *attributionist* proposals to explain why culpability is appropriate with respect to negligent agents of the sort that Randy illustrates (Arpaly 2003; Smith 2005; Sripada 2016). According to them, judging these agents as culpable is warranted because their unwitting wrongdoings reveal a moral inappropriate valuation structure. They care too much about things that they *should* not care. Or they care too little about things of considerable *moral importance*.

Attribution is obviously close cousin of voluntarism. It is, at the very least, an attempt to vindicate the thesis that culpability is grounded on a guilty mind, in the light of the possibility of inadvertent wrongdoing. The agent might not knowingly do something wrong. But the motivations on which she acts reflect something moral inappropriate about her structure of care.

The problem, of course, is that attributions of this kind are of limited applicability. No doubt, there are cases in which people fail to do wrong without even noticing it because they do not care enough or they do not care about these

things. The first interpretation of Randy's case is an example. But, clearly, what Randy's story under the second interpretation shows is that this is not always the case. In general, slips of this sort are evidence that the motivational force of an attitude and its power to influence one's conduct at a time can come apart.¹⁸ Inferring a guilty mind from the person's conduct in these cases is, therefore, simply a bad inference.

7. Skepticism

Negligence is a *morally significant* phenomenon. Or so I have argued up to this point. On the one hand, it is a distinctive fact about a portion of our psychology and our social life that some forms of wrongdoing are conceptualized as negligent and that negligent actors are judged as culpable. On the other hand, accounting for culpable negligence requires discussing ideas that run deep through common ways of theorizing about moral responsibility, ultimately casting doubt on the possibility a unified account of the phenomenon.

¹⁸ It is often the case that these two aspects of our motivational states are treated as one and the same thing. Many actions theorists, for instance, hold the view that behavior, as economists would put it, "reveal preference": one's preference function is manifested in behavior (for discussions and criticisms of the view, see Amaya 2013). Some have thought that behind this, there is a psychological law: the so-called "law of desire" (see Clarke 1994 for discussion and examples). There are, however, also many theoretical models where the strength of a motivation and its power to influence behavior are clearly kept apart (for discussion of some of these models, see Schroeder 2015).

These two points, as mentioned at the outset, are intimately connected. The theoretical importance of negligence is, no doubt, underscored by its practical significance. Thus, at this point one might reasonably wonder whether the practice of holding people responsible for their negligent behavior can be legitimized at all. Is the lack of an overarching plausible theory a sign that some of our practices need to be revised?¹⁹

It is possible, of course, that the lay and the experts got things wrong. Perhaps responsibility theorists have some *prima facie* methodological pressure to be conservative with respect to common intuitions and widespread practices. But those reasons can, obviously, be overridden. For instance, if these intuitions and practices are inconsistent with each other, it is clear that not all them can be coherently vindicated. Likewise, independent reasons for thinking that these intuitions and practices are the product of systematic error can make them lose their evidential credentials (for a thorough discussion of these issues, see Vargas 2013, part 1).

In line with this last point, some theorists have expressed skepticism about culpability for negligence (see, for instance, Levy 2017; Talbert 2017). Their hypothesis is that attributions of culpability based on it are the product of some systematic error in the form of a *problematic inference*. In brief, people witness an agent behaving in an inappropriate way or her action having morally objectionable

¹⁹ Moore and Hurd (2011: 192), who are skeptics of negligence as a category of legal culpability, develop this line. In their opinion, negligence is something of a “dog’s breakfast”: a residual category whose members have nothing in common except that they do not belong to other categories.

consequences. And they infer from it that the agent acted with some reprehensible motive or out of some censurable attitude.

This is an interesting hypothesis. If these theorists are right, negligence judgments result from an inappropriate inference from outcome to intent, in a way that make them highly-charged, psychologically speaking. More significantly, however, the hypothesis would provide an indirect vindication of the descriptive adequacy of the intuitions that initially recommended voluntarism as a theory of responsibility. People would make the problematic inference because, at least implicitly, they think moral responsibility is a matter of intent.

Although this hypothesis, if true, would go a long way towards recommending skepticism, it is not clear that it is adequately supported. There is evidence, as the skeptics of negligence point out, that human subjects tend to engage in reasoning about mental states when they are asked to produce judgments in response to morally charged actions and given no psychological information (for a review see Young and Tsoi 2013). But this evidence is hardly decisive here. Reasoning about mental states could be a matter of reasoning about the mental state that the person should have had, not about those that she actually has.

In fact, most of the evidence we have runs in the direction opposite to the hypothesis. Subjects reliably describe some behavior as negligent and attribute culpability to their actors before learning whether their behavior had undesirable consequences (Karlovac and Darley 1998). They hold on to judgments of culpability for negligent wrongdoing in the face of affirming that the actors didn't know or believe they were doing wrong (Nuñez et al 2014, study 3; Kneer and Machery 2018, study 4a). Indeed, their attributions of culpability are unaffected by

their explicit knowledge of absence of a desire to hurt or the presence an adequate attitude of care (Laurent et al 2016, experiment 2; Murray et al 2018).

This is not to say that we should take common intuitions about concrete instances of negligence at face value. As we have seen, foreseeability judgments tend to be afflicted by the so-called “hindsight bias.” And some forms of assessment of negligent agents of the sort that attributionists recommend might be signs of a larger attribution error. None of this, however, shows what the skeptic needs as evidence to undermine the practices described here as *products of a systematic error*. We do make some systematic errors in attributing culpability to concrete negligent actors. But that falls short of showing that there is a systematic error in taking negligence as adequate grounds for establishing culpability.

8. Caution

Throughout this chapter, I have emphasized how pervasive our practices are of holding people responsible for negligent wrongdoing. But it is important not to overstate this point. Ideally, we want a theory that account for our practices. But, in reality, there are significant portions of our practices that we do not know well enough, let alone understand.

As we have seen in the examples discussed here, concrete episodes of negligence can be interpreted in a variety of ways. So, whereas it is significant to note that people regularly find negligent agents culpable and even deserving of punishment, it is not entirely clear why they do so. This is, to be sure, a place where more work needs to be done. We need a diet of examples and scenarios that does a better job isolating the varieties of negligent wrongdoing. And we need to become more aware that even one scenario can be analyzed in multiple ways.

Another thing not to overstate is the level of existing agreement. With respect to some episodes of negligence, anecdotally or by scientific probing of lay attitudes, we can be confident that there is a robust agreement regarding the culpability of the agent. But there are plenty of real life cases where people, lay and experts alike, do not seem to agree. And, for the most part, we have little idea how big is the variation and what accounts for it.

As an illustration, consider a case that truly shows why discussions of negligence are morally significant. Every summer in the US approximately 30 toddlers die of hyperthermia after being *inadvertently* left in the backseat of the car (Weingarten 2009; Amaya 2013). The stories are sadly similar. A parent leaves home planning to drop his child at day care on the way to work. But instead he winds up driving straight to the office, leaving the child in the hot car all day.

Despite being in the public eye since they increased in frequency after new child-seat laws were enacted in 1998, these slips keep on occurring with regularity every year. Importantly for present purposes, their puzzling incorrigibility correlates with wide variation in perceptions of responsibility. Public reactions to them range from expressions of anger to sentiments of compassion and pity towards the parents. Legal precedence is all over the place. Some parents are never prosecuted; sentences for those who are prosecuted vary from one-year probation to 15 years in prison (Collins 2006).

No doubt, these episodes are real-life tragedies in need of a better explanation. We need to go beyond our usual stories about parents being too busy, too narrowly focused by work, etc (Amaya 2013, Amaya and Doris 2015). As we come to understand them more thoroughly, we also need to develop a better sense of what grounds the various responses that they generate. Perhaps people blame

the parents because they think they didn't care enough about their children. Maybe the tendency to exculpate them is explained by the thought that they have already suffered too much. At present we have some hypothesis. But we do not really have enough evidence to confirm them.

9. Conclusion

In this chapter, I have argued that negligence is a morally complex. It plays an important role in shaping up familiar ways in which moral evaluation takes place. It challenges widespread assumptions about the domain and the grounds for moral culpability. Trying to account for it, some promising avenues of moral theorizing are shown to be more problematic than they otherwise seemed.

In way, this is not an uplifting result. As it is has becomes apparent, we do not have yet a fully developed *moral psychology of negligence*. Concerning many of the core issues raised by it, there is more disagreement than consensus. Solving the disagreement often requires more evidence than we have in hand. But looking at the elements of the problem gathered here, does serve to show that developing a moral psychology of it is something worth doing.²⁰

²⁰ Research for this chapter was made possible by a generous grant (#60845) from the John Templeton Foundation to support the project "Getting Better at Simple Things: understanding and improving vigilant control." The views expressed in this chapter do not necessarily reflect the views of the Foundation. I would like to thank Fernando Rudy-Hiller, Samuel Murray, and Manuel Vargas for comments to an early draft of the chapter.

References

- Agule, C. (2016). Resisting Tracing's Siren Song. *Journal of Ethics and Social Philosophy* 10 (1):1-24.
- Alexander, L. & Ferzan, K. K. (2009). Against Negligence Liability. P. Robinson, S. Garvey, & K. K. Ferzan (Eds.), *Criminal Law Conversations* (Ch. 13). Oxford University Press.
- Alicke, M. D. (2014). Evaluating blame hypotheses. *Psychological Inquiry* 25 (2): 187-192.
- Amaya, S. & Doris, J. M. (2015). No Excuses: Performance Mistakes in Morality. In J. Clausen & N. Levy (Eds.), *Handbook of Neuroethics*. Springer. pp. 253-272.
- Amaya, S. & Murray, S. (ms). Vigilance.
- Amaya, S. (2013). Slips. *Noûs* 47(3): 559–576.
- Arpaly, N. (2003). *Unprincipled Virtue: An Inquiry into Moral Agency*. New York: Oxford University Press.
- Baron, M. (2001). I Thought She Consented. *Philosophical Issues*, 11: 1-32.
- Brink, D. O. & Nelkin, D. K. (2013). Fairness and the Architecture of Responsibility. In D. Shoemaker (Ed.), *Oxford Studies in Agency and Responsibility* 1 (pp.284-313). Oxford University Press.
- Chomsky, N. (1964). *Current Issues in Linguistic Theory*. The Hague: Mouton.
- Chomsky, N. (1965). *Aspects of the Theory of Syntax*. Cambridge: M.I.T. Press.
- Clarke, R. (1994). Doing what one wants less: a reappraisal of the law of desire. *Pacific Philosophical Quarterly* 75: 1–10.
- Clarke, R. (2014). *Omissions: Agency, Metaphysics, and Responsibility*. Oxford University Press.
- Collins, J. M. (2006). Crime and parenthood: The uneasy case for prosecuting negligent parents. *Northwestern University Law Review* 100 (2): 807-855.
- Cushman, F. (2008). Crime and punishment: Distinguishing the roles of causal and intentional analyses in moral judgment. *Cognition* 108 (2):353-380.
- Cushman, F., Sheketoff, R., Wharton, S. & Carey, S. (2013). The development of intent-based moral judgment. *Cognition* 127 (1): 6.
- Duff, R. (1993). Choice, Character, and Criminal Liability. *Law and Philosophy* 12 (4): 345-383.

- Feinberg, J. (1987). *Harm to Others*. New York: Oxford University Press.
- Fincham, F. & Jaspars, J. M. (1980). Attribution of Responsibility: From Man the Scientist to Man As Lawyer. In *Advances in Experimental Social Psychology* 13: 81-138.
- Finkelstein, C. (2005). Responsibility for Unintended Consequences. *Ohio State Journal of Criminal Law* 2 (2): 579-599.
- Fischer, J. M. & Ravizza, M. (1994). *Perspectives on Moral Responsibility*. Ithaca: Cornell University Press.
- Fischer, J. M. & Ravizza, M. (1998). *Responsibility and Control: A Theory of Moral Responsibility*. Cambridge University Press.
- Fischer, J. M. & Tognazzini, N. A. (2009). The truth about tracing. *Noûs* 43 (3): 531-556.
- Fischhoff, B. (1975). Hindsight is not equal to foresight: The effect of outcome knowledge on judgment under uncertainty. *Journal of Experimental Psychology: Human Perception and Performance* 1(3): 288-299.
- Fitzpatrick, W. J. (2008). Moral Responsibility and Normative Ignorance: Answering a New Skeptical Challenge. *Ethics* 118: 589-613.
- Graham, P. (2017) The Epistemic Condition on Moral Blameworthiness: a Theoretical Epiphenomenon. In P. Robichaud & J. W. Wieland, *Responsibility - The Epistemic Condition* (pp. 163-179). Oxford University Press.
- Guerrero, A. A. (2017). Intellectual Difficulty and Moral Responsibility. In P. Robichaud & J. W. Wieland, *Responsibility - The Epistemic Condition* (pp. 199-218). Oxford University Press.
- Haji, I. (1997). An Epistemic Dimension of Blameworthiness. *Philosophy and Phenomenological Research* 57 (3): 523 - 544.
- Hamilton. (1978). Who is responsible? Toward a social psychology of responsibility attribution. *Social Psychology* 41: 316-328
- Harman, E. (2011). Does moral ignorance exculpate? *Ratio* 24 (4):443-468.
- Helwig, C. C., Zelazo, P. D., & Wilson, M. (2001). Children's judgements of psychological harm in normal and noncanonical situations. *Child Development* 72: 66-81.
- Howarth, D (2006) Many Duties of Care: Or a Duty of Care? *Oxford Journal of Legal Studies*, 26(3): 449-472.

- Husak, D. (2011). Negligence, Belief, Blame and Criminal Liability: The Special Case of Forgetting. *Criminal Law and Philosophy* 5 (2):199-218.
- Kamin, K. A. & Rachlinski, J. J. (1995). Ex Post \neq Ex Ante: Determining Liability in Hindsight. *Cornell Law Faculty Publications* 646.
- Karlovac, M. & Darley, J. M. (1988). Attribution of Responsibility for Accidents: A Negligence Law Analogy. *Social Cognition* 6 (4): 287-318.
- King, M. (2017). Against Personifying the Reasonable Person. *Criminal Law and Philosophy* 11 (4):725-732.
- Kneer, M. & Machery, E. (2019). No luck for moral luck. *Cognition* 182: 331-348.
- Kohlberg, L. (1969). Stage and sequence: The cognitive development approach to socialization. In D. A. Goslin (Ed.). *Handbook of socialization theory* (pp. 347-480). Chicago, IL: Rand McNally.
- LaBine, S. & LaBine, G. (1996). Determinations of Negligence and the Hindsight Bias. *Law and Human Behavior* 20 (5): 501-516.
- Lagnado, D. A. & Channon, S. (2008). Judgments of Cause and Blame: The influence of Intentionality and Foreseeability. *Cognition* 108: 754-770.
- Laurent, S. M., Nuñez, N. L. & Schweitzer, K. A. (2016) Unintended, but still blameworthy: the roles of awareness, desire, and anger in negligence, restitution, and punishment. *Cognition and Emotion* 30 (7): 1271-1288.
- Levy, N. (2009). Culpable Ignorance And Moral Responsibility: A Reply To FitzPatrick. *Ethics* 119 (2009): 729-741
- Levy, N. (2014). *Consciousness and Moral Responsibility*. Oxford University Press.
- Levy, N. (2017). Methodological Conservatism and the Epistemic Condition. In P. Robichaud & J. W. Wieland, *Responsibility - The Epistemic Condition* (pp. 252-265). Oxford University Press.
- Margoni, F., Geipel, J., Hadjichristidis, C., & Surian, L. (2019). The Influence of Agent's Negligence in Shaping Younger and Older Adults' Moral Judgment. *Cognitive Development*, 49, 116-126.
- McBride, N. (2004). Duties of Care: Do they Really Exist? *Oxford Journal of Legal Studies*, 24(3): 417-441.
- McKenna, M. (2011). *Conversation and Responsibility*. Oxford University Press USA.

- Moore, M. & Hurd, H. (2011). Punishing the Awkward, the Stupid, the Weak, and the Selfish: The Culpability of Negligence. *Criminal Law and Philosophy* 5 (2): 147-198.
- Moore, M. (1997). *Placing Blame: A Theory of the Criminal Law*. Oxford University Press.
- Moore, M. S. & Hurd, H. M. (2002). Negligence in the Air. *Theoretical Inquiries in Law* 3 (2).
- Murray, S. (2017). Responsibility and vigilance. *Philosophical Studies* 174 (2): 507-527.
- Murray, S., Murray, E. D., Stewart, G., Sinnott-Armstrong, W. & De Brigard, F. (2018). Responsibility for forgetting. *Philosophical Studies*: 1-25.
- Nelkin, D. K. & Rickless, S. C. (2017). Moral Responsibility for Unwitting Omissions: A New Tracing View. In *The Ethics and Law of Omissions* (pp. 106-129). New York, NY, USA.
- Nobes, G., Panagiotaki, G. & Bartholomew, K. J. (2016). The influence of intention, outcome and question-wording on children's and adults' moral judgments. *Cognition* 157: 190-204.
- Nobes, G., Panagiotaki, G. & Engelhardt, P. (2017). The development of intention-based morality: The influence of intention salience and recency, negligence, and outcome on children's and adults' judgments. *Developmental Psychology* 53 (10): 1895-1911.
- Norman, D. (1981). Categorization of action slips. *Psychological Review* 88: 1-15.
- Núñez, N., Laurent, S. & Gray, J. M. (2014). Is Negligence a First Cousin to Intentionality? Lay Conceptions of Negligence and Its Relationship to Intentionality. *Applied Cognitive Psychology* 28: 55-65.
- Owen, D. G. (2007). The Five Elements of Negligence. *Hofstra Law Review* 35 (4): 1671-1686.
- Pereboom, D. (2015). A Notion of Moral Responsibility Immune to the Threat from Causal Determination. In R. Clarke, M. McKenna & A. M. Smith (eds.), *The Nature of Moral Responsibility* (pp. 281-296). Oxford University Press.
- Piaget, J. (1932). *The Moral Judgment of the Child*. New York: Free Press.
- Raz, J. (2010). Responsibility and the Negligence Standard. *Oxford Journal of Legal Studies* 30 (1): 1-18.

- Reason, J. (1984). Lapses of attention in everyday life. In R. Parasuraman & D. Davies (Eds.), *Varieties of attention* (pp. 515–549). New York: Academic Press.
- Robichaud, P. & Wieland, J. W. (eds.) (2017). *Responsibility - The Epistemic Condition*. Oxford University Press.
- Roese, N. J. & Vohs, K. D. (2012). Hindsight Bias. *Perspectives on Psychological Science* 7: 411-426.
- Rosen, G. (2004). Skepticism about moral responsibility. *Philosophical Perspectives* 18 (1): 295–313.
- Rudy Hiller, F. (2017). A Capacitarian Account of Culpable Ignorance. *Pacific Philosophical Quarterly* 98 (S1):398-426.
- Rudy-Hiller, F. (Fall 2018 Edition). The Epistemic Condition for Moral Responsibility. *The Stanford Encyclopedia of Philosophy*, Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/fall2018/entries/moral-responsibility-epistemic/>>.
- Schroeder, T. (Summer 2017 Edition). Desire. *The Stanford Encyclopedia of Philosophy*, Edward N. Zalta (ed.), URL = <<https://plato.stanford.edu/archives/sum2017/entries/desire/>>.
- Shabo, S. (2015). More Trouble with Tracing. *Erkenntnis* 80 (5): 987-1011.
- Sher, G. (2006). Out of control. *Ethics* 116 (2): 285-301.
- Sher, G. (2009). *Who Knew?: Responsibility Without Awareness*. Oxford University Press USA.
- Shiffrin, S. (2017). The Moral Neglect of Negligence. In D. Sobel, P. Vallentyne, and S. Wall (Eds.), *Oxford Studies in Political Philosophy, Volume 3*. Oxford University Press.
- Shultz, T. R., & Wright, K. (1985). Concepts of negligence and intention in the assignment of moral responsibility. *Canadian Journal of Behavioural Science* 17 (2): 97-108.
- Simons, K. (2009). The distinction between negligence and recklessness is unstable. In P. Robinson, K. Ferzan & S. Garvey (Eds.), *Criminal Law Conversations* (pp. 290—291). Oxford University Press.
- Simons, K. (2011). When is Negligent Inadvertence Culpable? *Criminal Law and Philosophy* 5 (2): 97-114.

- Smith, A. (2005). Responsibility for Attitudes: Activity and Passivity in Mental Life. *Ethics* 115: 236–71.
- Smith, H. (2011). Non-Tracing Cases of Culpable Ignorance. *Criminal Law and Philosophy* 5 (2): 115-146.
- Sripada, C. (2015). Moral Responsibility, Reasons, and the Self. In D. Shoemaker (Ed.), *Oxford Studies in Agency and Responsibility: Volume 3*. Oxford University Press.
- Talbert, M. (2013). Unwitting Wrongdoers and the Role of Moral Disagreement in Blame. In D. Shoemaker (ed.), *Oxford Studies in Agency and Responsibility: Volume 1*. Oxford University Press.
- Talbert, M. (2017). Omission and Attribution Error. In D. K. Nelkin & S. C. Rickless (Eds.), *The Ethics and Law of Omissions* (pp. 17-35). Oxford University Press.
- Tappolet, C. (2004). Through thick and thin: Good and its determinates. *Dialectica* 58 (2): 207-220.
- Timpe, K. (2011). Tracing and the Epistemic Condition on Moral Responsibility. *Modern Schoolman* 88 (1/2): 5-28.
- Vargas, M. (2005). The Trouble with Tracing. *Midwest Studies in Philosophy* 29 (1): 269-290.
- Vargas, M. (2013). *Building Better Beings: A Theory of Moral Responsibility*. Oxford University Press.
- Vargas, M. (Forthcoming). Negligence and Social Self-Governance. In A. Mele (Ed.), *Surrounding Self Control*.
- Weingarten, G. 2009. Fatal distraction, *The Washington Post*, 03/08/09. Retrieved on Feb 1, 2010 from <http://www.washingtonpost.com/wpdyn/content/article/2009/02/27/AR2009022701549.html>
- Wieland, J. W. (2017). The Epistemic Condition. In P. Robichaud & J. W. Wieland, *Responsibility - The Epistemic Condition* (pp. 1-28). Oxford University Press.
- Wolff, S. (1990). *Freedom within reason*. New York: Oxford University Press.
- Young, L. & Tsoi, L. (2013). When Mental States Matter, When They Don't, and What That Means for Morality. *Social and Personality Psychology Compass* 7/8: 585-604.

- Young, L., Cushman, F., Hauser, M. & Saxe R. (2007). The neural basis of the interaction between theory of mind and moral judgment. *Proceedings of the National Academy of Sciences of the United States of America* 104 (20): 8235-40.
- Young, L., Nichols, S. & Saxe, R. (2010). Investigating the Neural and Cognitive Basis of Moral Luck. *Review of Philosophy and Psychology* 1 (3): 333-349.
- Zelazo, P.D., Frye, D. & Rapus, T. (1996). An age-related dissociation between knowing rules and using them. *Cognitive Development* 11: 37-63.
- Zimmerman, M. J. (1986). Negligence and moral responsibility. *Noûs* 20 (2): 199-218.
- Zimmerman, M. J. (1997). Moral responsibility and ignorance. *Ethics* 107 (3): 410-426.
- Zimmerman, Michael J. (1988). *An Essay on Moral Responsibility*. Rowman & Littlefield.