



Author's Reply: Negligence and Normative Import

Katrina L. Sifferd¹ · Tyler K. Fagan¹

Accepted: 22 July 2021

© The Author(s), under exclusive licence to Springer Nature B.V. 2021

Abstract

In this paper we attempt to reply to the thoughtful comments made on our book, *Responsible Brains*, by a stellar group of scholars. Our reply focuses on two topics discussed in the commenting papers: first, the issue of responsibility for negligent behavior; and second, the broad claim that facts about brain function are normatively inert. In response to worries that our theory lacks normative implications, we will concentrate on an area where our theory has clear relevance to law and legal policy: juvenile responsibility.

1 Negligence

Several of our commenters focused on our handling of negligence, no doubt because negligence is where the rubber meets the road in a theory of legal responsibility. And as Dennis Patterson notes in his comment, we feel our handling of negligence cases shows our theory to be superior to some other recent scientific approaches to responsibility (Levy 2014).

The discussion of culpable negligence in *Responsible Brains* (hereafter, RB) uses the real-life example of Bert (not his real name), who was arrested for child neglect in Illinois after appearing to honestly forget that it was his weekend to watch his kids (133). Instead of meeting them at his apartment, Bert went straight from work to the airport to fly to Las Vegas, leaving his two young children—who were “latchkey kids”—home alone for an entire weekend. Under Illinois law, Bert is only responsible for child abandonment if he “knowingly” left his children unattended. Yet the judge in his case found him guilty, in a result that coheres with the law in other states and, arguably, our commonsense judgments regarding these sorts of cases. “The state doesn’t just let parents forget about their children,” announced the judge as she handed down her verdict.

✉ Tyler K. Fagan
tyler.fagan@elmhurst.edu

¹ Department of Philosophy, Elmhurst University, 190 Prospect Avenue, Elmhurst, IL 60126, USA

When our forgettings violate an important standard of care, such as the one required of parents for their children, and this violation causes harm, blame seems appropriate. But what grounds blame (and possibly punishment) in such cases? On most theories of responsibility, we are justified in blaming a person for causing harm if that harm was intentional; or, if the person knew of a risk of harm and performed the risky behavior anyway. Similarly, criminal guilt typically requires both an act that causes criminal harm (*actus reus*) and that a defendant possessed a culpable mental state (*mens rea*) related to that criminal harm. So, a person might be found guilty of a homicide if they acted for the *purpose* of causing a death or *knowing that* a death might result.

Lapses or forgettings may be prosecuted as negligence, however, even where there is no conscious desire to cause harm, nor any belief that harm may result from an act. In negligence cases there is not even a conscious disregard of the risk of harm; for this constitutes recklessness, and if the forgetting is honest, one did not know there was a risk of harm at the time of the harmful action. Other cases of negligence include circumstances where a person caused themselves not to be able to form the requisite *mens rea*, and then caused harm (some such cases involve voluntary intoxication). As we will see, lapse and intoxication cases are quite different, even though both have been prosecuted as criminal negligence. On first glance, one might think that what unifies these cases is that negligent persons are responsible for the harm they caused because, even in the absence of an intentional or knowing mental state, they *ought to have acted differently*—they *should have paid better attention, should not have driven home*, etc. But as we will see, things are more complicated than this.

On the particular theory of responsibility articulated in RB, a person is responsible only if they possess a Minimal Working Set of executive functions (MWS). These brain functions are necessary to what philosophers call “reasons-responsiveness” (Fischer and Ravizza 1998) and what H.L.A. Hart termed “capacity responsibility,” which he claimed was a precondition for legal liability (1968). A person who is responsive to moral and legal reasons has the capacity to act in accordance with those reasons via the executive functions of attentional focus, task-shifting, planning, and inhibition. Persons who are reasons-responsive are responsible when they act because they have what Fischer and Ravizza call “guidance control” over that act, which justifies application of blame and, in some cases, criminal guilt and punishment. This is our theory of responsibility in a nutshell.¹

A person who is negligent fails to be responsive to an important moral or legal reason related to the harm they cause. The primary question with regard to negligence, then, is whether agents are responsible for such failures. Hart understood culpable negligence as a failure to exercise capacity responsibility, specifically “the

¹ Many of our critics wonder what we are offering in addition to Hart’s and Fischer and Ravizza’s theories of responsibility. That is, surely it is interesting to speculate about the function and even location of brain activity related to responsible action; but does this information do any normative work beyond what is already implied in a “standard” reasons-responsive theory? In a later section of this response we will try to show why we think the answer is “yes.”

capacity to advert to, and think about and control, conduct and its risks” (1968, p. 157). Others argue that failure alone is not enough to ground responsibility; we need some additional inculpatory factor to ground negligent responsibility (Stark 2016). This inculpatory factor might be some synchronic or diachronic state of affairs, often having to do with an agent’s rational or cognitive system.

In RB we claim that negligent persons like Bert are indeed responsible for their acts, although less responsible than persons who intentionally cause harm. Many of our commentators agree with us that persons can be criminally responsible for negligent behavior, but take issue with our analysis of the grounds for negligent responsibility; these commentators include Moore, Agule, and Patterson. Douglas Husak, however, disagrees that negligent actors are responsible at all. We will begin with his concerns.

a. Husak on RB on Negligence

Husak worries that the neuroscience we utilize in RB neither informs nor challenges the “established orthodoxy in the Anglo-American legal tradition” on the topic of negligence (Husak this volume). He also claims our account of negligence “misrepresents” that traditional orthodoxy. Husak seems to indicate that we shouldn’t assume criminal negligence is something we must account for, since “perhaps most” legal philosophers reject it as indefensible; at any rate, he argues that our defense of it is inadequate. We won’t attempt to sort out how many legal philosophers and scholars doubt that people should (sometimes) be held responsible on the basis of negligence; it is enough to note that Husak himself is in this group; although others are not (see, for example, Stark 2016; Sarch, 2018; Baron 2020; Moore and Hurd 2011), and that negligence is still a category of *mens rea* under the Model Penal Code (MPC), including Section 210.4, which lists negligent homicide as a felony.

As indicated above, attempts to defend culpable negligence include identifying some state of a person’s rational system as the inculpatory factor causally related to the harm. Such approaches may involve “tracing” to seek an inculpatory factor in the person’s past. Another tactic is using counterfactuals to describe conditions under which the negligent party failed to attend risks when they should have. Our aim in RB was to reject tracing and offer a diachronic inculpatory factor, although it may appear that we waffle a bit between offering an inculpatory factor and looking to counterfactuals.

One example of the “inculpatory factor” approach is Garvey (2006), who argues that negligent responsibility is justified where a person has doxastic self-control. Negligent culpability, says Garvey, consists not in any prior actual choice to do wrong, or imagined hypothetical choice to do wrong, but in the culpable failure to exercise control over one’s beliefs. Stark (2016) expands this theory, including within the sphere of doxastic self-control the ability to manage one’s beliefs and character traits. Stark then claims we can hold persons responsible when they have background beliefs that could have led to an understanding of the

risk, even where the belief regarding the risk is not actually formed (2016, pp. 247–249).

In RB we do something similar: we ground responsibility for unknown risks in the diachronic possession of an MWS, which we claim enables the doxastic self-control discussed by Stark. More generally, executive functions enable the capacity for self-formation and self-governance (Roskies 2018); and we claim persons with an MWS can manage not only their beliefs but their own reasons-responsiveness, developing dispositions to be sensitive to moral and legal reasons, and to be norm-abiding. For this reason, we are comfortable claiming that a person can be responsible for a risk even if their MWS was not engaged with regard to that risk. A person with diachronic possession of an MWS is responsible if they *could have engaged their MWS* in such a way that the risk was acknowledged, even if they did not. MWS is our inculpatory factor, but we offer a counterfactual explanation as to why the MWS grounds responsibility in such cases.

Husak claims that to ground responsibility for negligent acts, our position requires that agents be responsive to objective—rather than subjective—reasons for acting. By his lights, our view demands that agents be responsive to reasons of which they are not, and could not reasonably be expected to be, aware at the time of their negligent action. This demand “clearly exceeds the capacities of any human brain,” Husak argues; all one can do “at a given time...is to respond to the reasons of which [one is] aware” (5).

We wonder who is bucking the legal orthodoxy here. As other commenters have noted, criminal negligence persists as a legal doctrine, and under negligence laws persons are held responsible when they *ought to have known their actions created a substantial risk*, where that risk causes criminal harm. The criminal law does indeed hold people responsible for things they were not—but ought to have been—aware of. And we think this is correct because of the diachronic capacities discussed above. Husak may think the law should not demand this much of an agent, but that complaint does not seem an objection to our theory of responsibility in particular.

One of us has expressed some worries about Husak’s theory of criminal responsibility elsewhere (Sifferd 2018). Husak (2016) adopts a unique *subjective* version of Fischer and Ravizza’s (1998) reasons-responsive theory of responsibility. Judgments of moral responsibility, he says, require an “internal assessment of the reason-responsiveness of agents” (Husak 2016, p. 154).

But Fischer and Ravizza’s theory revolves around the idea that we are morally responsible for an action when the mechanism that issues that action is *moderately* reasons-responsive in an appropriate way. To be reasons-responsive an agent must have been regularly receptive to a range of reasons such that she manifests an intelligible pattern of responsiveness (over time) from a third person perspective (Fischer and Ravizza 1998, p. 71). In addition, it must be shown an agent would react to at least one sufficient reason to do otherwise than she did in some possible scenario—that is, she might have acted differently in some close possible world (44). However, it does not follow that the agent could have responded differently to the *actual* reasons. To be responsible, Fischer and Ravizza require only that a person possess a mechanism that is, over time, moderately sensitive to reasons of the right type;

and again, they offer counterfactual explanations as to why this sort of mechanism grounds responsibility.

Of course, there are many kinds of negligence cases. In those where someone was distracted from a risk, or created a risk due to a lapse, it seems straightforward that the possession of an MWS and guidance control can ground responsibility, whether one frames this as an inculpatory factor or prefers the counterfactual analysis that one *could have* responded to the relevant reasons. But as Husak notes, there are other negligence cases where persons seemingly do not possess an MWS at the time of action, including voluntary intoxication. In these cases, one might wonder how executive functions can be thought to be involved in the action related to the risk of harm at all (5). We discuss such cases when handling criticisms from other commenters, below.

Finally, Husak wonders about how our theory handles cases of factual mistakes, given we are willing to hold persons responsible who do not know they are performing risky behavior. Husak claims we are only responsible for what we know; even if we have stupid factual or moral beliefs that are mistaken such that a reasonable person wouldn't hold them, we aren't responsible when we harm others based on a mistaken understanding regarding risk. This claim is directly related to his belief that the sole purpose of criminal law and punishment is to address moral wrongs by delivering just deserts. One who acts without knowing the risks isn't morally blameworthy.

In RB, we support a hybrid justification for criminal law and punishment. One of us has written elsewhere about the fit between forward-looking and backward-looking justifications for criminal law (Sifferd, forthcoming). Hart (1968) appeals to consequentialist principles as the “general justifying aim” of criminal punishment, although in practice forward-looking factors seem to play a less important role than retribution in justifying criminal punishment, especially if one looks to legislative codes and documents.² In the US, the state typically claims hybrid justification (both forward- and backward-looking) when assigning punishments, where retribution tends to be the primary aim, with deterrence considered within the parameters established by considerations of “just deserts”.

Forward-looking justifications other than deterrence, however, can be offered for our punishment practices: (1) the criminal law makes society's moral rules more salient, and in this way encourages compliance with those rules; and (2) holding persons criminally responsible has the forward-looking effect of assigning, supporting, and even enhancing the moral agency that underpins responsible action. We

² For example, members of the American Law Institute (ALI) voted to approve a new final draft of Model Penal Code: Sentencing in 2017. In the newest iteration of the MPC, which is meant as a guiding document for all state penal codes, the stated purpose of criminal sentencing is “to render sentences in all cases within a range of severity proportionate to the gravity of offenses, the harms done to crime victims, and the blameworthiness of offenders.” When “reasonably feasible” the state should also attempt “to achieve offender rehabilitation, general deterrence, incapacitation of dangerous offenders, restitution to crime victims, preservation of families, and reintegration of offenders into the law-abiding community, provided these goals are pursued within the boundaries of proportionality...” Sect. 1.02 Model Penal Code: Sentencing (Am. Law Inst., Proposed Final Draft, 2017).

ground these claims in social theories of responsibility, such as Victoria McGeer's *scaffolded responsibility* (McGeer 2015; McGeer and Pettit 2013) and Manuel Vargas's *agency cultivation* (Vargas 2013), which emphasize the ways moral agency can be scaffolded upon and bolstered by cultural norms and a societal audience.

Interestingly, mistakes of fact and mistakes of law—which often boil down to mistakes about what is right and wrong—come apart when we consider forward-looking aims of criminal law. In part due to the relative accessibility of information about moral and legal norms, mistakes of law seem to be within a person's control in a way that mistakes of fact usually aren't, which means that persons can be incentivized by one's societal audience and the law to avoid them. Mistakes of fact, on the other hand, are often the sort that a reasonable person would make under the circumstances. We think people ought to be excused for reasonable mistakes of fact—but not law—in part because the threat of punishment is unlikely to make persons less likely to make a reasonable mistake with regard to their perceptions. It seems unfair for the criminal law to demand that people be more than reasonably careful with regard to their factual understandings. It is reasonable, however, for the threat of punishment to encourage awareness of moral and legal norms. Someone considering using a tax haven may be disposed to investigate the fine points of tax law; gun owners may be motivated to investigate state castle laws. In the same vein, it is fair to ask people with diachronic possession of an MWS to exercise guidance control and reasonable care so as not to create risks that cause criminal harm.

b. Moore on RB on Negligence

Another commentator, the estimable Michael S. Moore, says he agrees with the broader strokes of our argument but claims our theory does not properly manage negligence cases. At the root of Moore's critique is a fairly fundamental worry: that we have adopted the wrong version of compatibilism—namely, reasons-responsive compatibilism—and, to make matters worse, have amended the standard reasons-responsive approach in unhelpful ways. Before returning to the specific topic of negligence, we will address this more general line of criticism.

Moore takes us to task for not giving our readers a full taxonomy of compatibilist theories (an entirely fair charge). He distinguishes *source* compatibilism and *conditionalist* compatibilism as the two major camps of compatibilist thought. The views we emphasize in RB, namely reasons-responsive and deep-self theories, are both species of source compatibilism. Conditionalist compatibilism, which we do not discuss in RB, requires for responsible action that a person could have chosen other than she did (adopting the principle of alternate possibilities, or PAP).

Moore alleges that in RB, despite not acknowledging the gulf between source and conditionalist compatibilism, we betray our awareness of that divide and align ourselves with one side of it, “put[ting] aside conditionalist analyses with [our] blunt conclusion that ‘the ability to do otherwise is not necessary to responsible action’” (Moore, this volume). This allegation rests on a misreading of the passage in question (one we may have invited). In the line quoted by Moore we are not asserting our own view, but rehearsing Harry Frankfurt's classic argument (1969) against the PAP.

The sentence begins “And if that is right [i.e., if agents are still responsible even in ‘Frankfurt cases’], then the ability to do otherwise is not necessary to responsible action” (Hirstein, Sifferd, and Fagan 2018, p. 48). Our position on whether responsible action requires the ability to do otherwise is not so simply put—and depends, we would say, on what “ability to do otherwise” means.

Putting this aside, there is an important substantive question raised by Moore’s critique: Have we chosen the wrong side by aligning with source compatibilism, and, as a result, have we left key concepts such as “capacity” under-analyzed?

The conditionalist camp, which Moore speaks for here, contends not that source compatibilism is utterly wrongheaded, but that it is *not enough* to ground responsibility. Moore claims that “responsibility for some wrongful action A requires more than that A be properly sourced within certain structures and capacities of the agent (although it does require at least that)” (6). The *more* he refers to here is a “freedom worth wanting:” in conditionalist terms, the freedom to do otherwise or the power of “choosing and doing one thing rather than another” (6). We suspect that the counterfactual freedom Moore embraces is quite close indeed to the freedom of an agent with guidance control over their actions. And if the divide between source and conditionalist compatibilists can be bridged in this way, maybe we can defuse Moore’s worry that source compatibilism is insufficient to the task of grounding responsibility.

On Moore’s brief restatement of the conditionalist view, responsibility for an action entails the ability to have chosen and done otherwise, analyzed in terms of what the agent *would* have done *if* certain other circumstances had obtained (Moore 2021, p. 7). Conditionalists of different flavors and vintages fill in the antecedent of this conditional with different details—Moore nods to both the “classic” and “new” subcamps of conditionalism—but the basic picture is the same. If one rejects libertarian contra-causal free will, what could this counterfactually specified ability come to? Moore has an entirely reasonable answer: “‘Ability,’ by this analysis, does not refer as much to facts about successful future behavior as it does to *inner dispositions and to the physical equipment that realizes them...*” (7, our emphasis).

It’s in this last move of conditionalist reasoning that we see the common ground shareable, if not explicitly shared, between conditionalist and source compatibilists. For conditionalists, responsibility for an act requires the counterfactually specified freedom to have chosen/done otherwise, which is purportedly *more* than just proper sourcing “within certain structures and capacities of the agent.” But possessing this all-important ability turns out to be mostly, if not entirely, a matter of having the right “inner dispositions” realized by an intact set of “physical equipment” (7). To our ears this sounds very much like the “structures and capacities” of the agent; it is not clear what the ability to do otherwise amounts to that would not be captured by a plausible source-compatibilist view of how responsible actions originate from within reasons-responsive mechanisms (or the deep self) of the agent.

Our point is not that Moore’s counterfactual language is only a roundabout way of describing something like an MWS, or that conditionalist compatibilism is “really” source compatibilism in disguise. Nonetheless Moore’s language seems consistent with many versions of source compatibilism, including Fischer and Ravizza’s reasons-responsive view and our (RB’s) elaboration of that approach;

indeed, it may be that each position understands its “primitives” in terms borrowed from the other. Consider Fischer and Ravizza’s (1998) guidance control: One of the criteria for the existence of even a weakly reasons-responsive mechanism is that the agent who possesses the mechanism *would react to at least one sufficient reason to do otherwise than the agent did in some possible scenario* (38).³ Thus a counterfactually specified ability is already implicit in this idea of a reasons-responsive mechanism, just as something like a reasons-responsive mechanism seems implicit in the idea of a counterfactual ability to do otherwise.

It is easy to imagine allegations of “collapse” being lodged from either side of the source-versus-conditionalist divide: conditionalists arguing that a reasons-responsive mechanism is ultimately a set of counterfactuals, and reasons-responsive theorists arguing that the relevant counterfactual “abilities” are ultimately picking out a reasons-responsive mechanism. It’s not clear who would have the upper hand in this debate or what would be illuminated by its playing out. Perhaps we ought to think of counterfactual abilities and reasons-responsive capacities as defined in terms of each other; perhaps one’s allegiance to source or conditionalist compatibilism depends on which side of this analysis one prefers to focus on. We think this is an interesting avenue for further exploration and we thank Moore for alerting us to it, but we don’t find it too troubling a prospect for our account of negligence in RB.

However, Moore’s critique includes other points worth addressing. Beyond his reservations about reasons-responsive views generally, he worries that our diachronic version of reasons-responsiveness is weaker than Fischer and Ravizza’s synchronic version. As indicated above, the diachronic component of our theory was meant to draw attention to the fact that moral and legal norms are a large part of the normative culture that human beings are asked to be sensitive to; and that these norms can interact with an MWS (or reasons-responsive mechanism) such that we can often shape our moral agency over time to be more risk-averse and law-abiding (Vargas 2013, 2015; McGeer 2015). As normative beliefs are developing, we think, dispositions to act in more or less norm-abiding ways are also under construction. We act not only in relation to synchronic mental states or decisions, but against a background of beliefs developed over time regarding the act’s normative quality (beliefs the agent could report if asked).

To return to the topic of negligence and the case of Bert the forgetful father: We claim that Bert is responsible for his harmful action because (1) he had an MWS at the time of the action *and* (2) he had an MWS throughout the time when his dispositions and habits regarding his parenting duties were being formed. Because of (1), Bert had guidance control with regard to the act that harmed his children even though he was unaware of the risk of harm (where under different conditions he might have been aware of the risk). And because of (2), he also had many *earlier* opportunities both to shape his disposition to be a mindful parent and to adjust for his shortcomings in that department. Moore claims our strategy of holding Bert

³ Another condition is that the agent must “own” their mechanism by taking responsibility for it. We note that one can do this by directing reactive attitudes toward one’s own actions; and in a diachronic sense, one might exhibit ownership by manipulating the operations of their mechanism over time.

responsible for his lapse based upon his possession of a diachronic MWS is a type of *tracing* account (see below), and vulnerable to an objection familiar to such views: If the mere possession of an MWS at the time of Bert's action is not enough to ground his responsibility, then any earlier failures we may blame him for—not keeping a calendar of his parenting commitments, for example—are likewise not his fault based upon mere possession of an MWS (2021, 9). And if the mere possession of an MWS at the time of Bert's action *is* enough to hold him responsible, then the diachronic MWS and its wider timescale seem not to be doing any work at all. This is an excellent point.

A few words on tracing accounts, by which some have attempted to ground responsibility in negligence cases via a technique where one traces back to a decision or action that led to the current harm for which one is now responsible (Fischer and Tognazzini 2009; Vargas 2005). Tracing looks into an agent's past in an attempt to find a culpable mental state that can be reasonably linked to the current harm. This technique, however, is problematic even in the cases for which it is most often used; namely, cases where an intoxicated agent causes harm. A drunk driver may lack the capacity to understand the risk of harm their driving poses, but it is typically extraordinarily difficult to identify a single decision or action that led to an incapacitating level of intoxication (Vargas 2005; King 2009). For which decision is the agent relevantly culpable: the decision to drive to the bar? To have the first drink (if they are an alcoholic), or the second, or the third?

We should have been clearer about the grounds for Bert's responsibility. He is responsible because he possessed an MWS both synchronically at the time of the act that caused harm and diachronically, even though he did not engage his MWS with regard to the risk that his kids might be left home alone. On our view, these are the “easy” negligence cases—where an actor unknowingly breaches an important standard of care and causes harm, but has both diachronic and synchronic possession of an MWS. Cases where there is only a diachronic, but not synchronic MWS (e.g. voluntary intoxication cases); or no diachronic, but a synchronic MWS (e.g. negligent swamp man) are significantly harder (and in the latter case, probably so rare as to be uninteresting). More on this below.

Moore is worried that our MWS/guidance control theory doesn't give an adequate account of ability or capacity, which is what we are pointing to as the diachronic/synchronic grounds of responsibility. On Moore's conditionalist view, Bert is culpable for not meeting his parenting obligations because he *could* have met them—meaning that, if the world were different (say we're in a world where Bert and his ex-partner were assigned a parenting coordinator that made them keep an online parenting calendar; or his ex-partner had called to remind him about the children's visit as he drove to the airport), Bert *would* have kept his obligations.

Aside from noting the fundamental compatibility between conditionalist and source-compatibilist accounts (see above), we are unsure that Moore's theory is more specific and useful than ours. Moore faces the difficult task of determining what features might be altered in an agent or their environment to indicate a range of ability. For example, in a world where Bert was taught to be organized by a fastidious parent, would he have kept his obligations? If yes, what does this mean regarding Bert's current ability to be a good parent? Or, in a world where a co-worker

asked Bert if he had his kids that weekend, would Bert have started to wonder about his schedule and asked his ex, keeping his obligation in front of his mind and meeting it? Questions like these are relevant to assessing the grounds for Bert's responsibility, and worth asking—but in our view they are probing, in their counterfactual way, whether Bert properly developed and possessed an MWS in the time leading up to his harmful action.

Moore admits, of course, that it is “enormously tricky” to frame the relevant counterfactuals needed to discern who had the ability to understand the risks associated with negligent behavior. But he says it is necessary to “distinguish two kinds of cases where some executive control function could have operated but didn't,” versus not operating because it couldn't have done. Again, it seems we might answer this question by noting that a person who possesses an MWS at the time of a crime, *and* in some period beforehand, certainly had the ability to have engaged their executive functions with regard to the risk—to be sensitive to it, to deliberate with regard to it, and to act or inhibit action based upon an understanding of it. To possess an MWS diachronically means that one's MWS engages in response to appropriate conditions (some persons with schizophrenia or REM behavioral disorder may be exemplars of a dysfunctional MWS).

Perhaps, however, Moore's view handles the intoxication case better than ours. We have a problem in such cases, because often the person acting doesn't have an MWS at the time of the act that caused harm—even if they did possess one in the time period running up to the risky behavior. On Moore's theory, though, it seems clear that the actual world could have changed in various simple ways so that someone who drove drunk and caused harm would not have done so. In such cases it seems we must hang responsibility on a person's possession of an MWS during the time frame before they were incapacitated alone; or indeed, engage in tracing, thereby holding the actor responsible for a particular decision or act performed earlier.

If culpable negligence always requires *both* synchronic and diachronic MWS, then a very intoxicated driver would not be culpable for negligence in hitting a child while drunk, provided (as seems probable) that he was temporarily without an MWS. This would seem to be an unacceptable result. Central to our theory—and to Hart's notion of capacity responsibility, and any capacitarian theory of legal responsibility, really—is that one is worthy of legal sanction only if the law can serve as a reason for action; or, only if one can understand legal and moral rules and conform one's behavior to those rules. In such cases we bite the bullet and opt to ground responsibility in negligence cases in a diachronic assessment of MWS, looking to see if one had an MWS over a period of time sufficient to review or otherwise affect one's decisions or acts with regard to the risk ultimately taken.

As already discussed, this emphasis on diachronic possession of an MWS acknowledges the ways in which identity and agency are built and maintained over time. The law does not provide reasons for or against acting only at the synchronic moment of decision or action; it is part of the normative milieu our reasons-responsiveness is scaffolded upon (McGeer and Pettit 2013; McGeer 2015), shaping our legal agency and motivating us to cultivate law-abiding dispositions.

Of course, responsibility for a negligent act cannot rest *solely* on diachronic possession of an MWS. If this were the case, then persons would be responsible for harm caused when they were *involuntarily* intoxicated. What makes a person responsible in a drunk driving case is that the person knew drinking and driving was a risky behavior that could cause harm; and because they possessed an MWS they had the ability to shape their mental states/dispositions/environment so as to avoid the risk of causing harm when intoxicated. We might be tempted to try to trace back to a particular moment when a person engaged their MWS with regard to those risks of drunk driving; but we believe (1) there are likely to be many such moments, and they are difficult to parse; and (2) it is equally important to responsibility that the person did not use their diachronic capacity for guidance control to avoid drunk driving. In other words, tracing would not provide us with a full explanation as to why a drunk driver is responsible.

c. Agule on RB on Negligence

In his intelligent and probing comments, Craig Agule (this volume) provides another example of negligence to test our theory: Jeri Dawn Montgomery was chatting on her cell phone when she realized that she was about to miss her entrance to the highway. She swerved to make the entrance, causing an accident that resulted in a death, and was convicted of criminally negligent homicide. Agule stipulates that Montgomery was unaware of the risk she was causing to others at the time she acted. Contra Moore (see below), Agule indicates that negligence cases like Montgomery's can be handled by theories like ours *without* looking to counterfactual conditions to ground responsibility, albeit with a few friendly amendments.

As Agule notes, the Montgomery case makes the difficulties with tracing back to a decision or action for which a person can be held responsible in cases of negligence clear. There does not seem to be an earlier moment at which Montgomery had knowledge regarding the risks involved in the behavior that caused the accident. There is "no particular point at which the agent was aware of the possibility of acting to avoid the later risk" (Agule, this volume). That is, there seems to be no moment where there is a culpable mental state (in terms of awareness of a risk of harm) to trace back to.

Like us, Agule is sympathetic to Hart's capacitarian strategy for handling negligence. That is, if one has the right sort of capacities necessary to legal agency at the time of a voluntary action, and they had a "fair opportunity to exercise those capacities," one has the "guidance control" necessary for legal responsibility. Much of intentional action flows from knowledge and desires that are in some sense unconscious, but this does not mean that these beliefs and desires are "not ours" or that we are not in control of behavior that issues from them; if that were true, no one would be responsible for anything. On our theory, a person has the right sort of control when their action is produced by a reasons-responsive mechanism, which underpins both the capacity to understand legal and moral rules and the ability to conform one's behavior to these rules.

Agule worries, however, that our theory is too permissive in classifying persons as responsible in cases of negligence because we hold persons responsible when a person's executive functions *should have* played a role in preventing an agent's act or omission (that is, the person possessed an MWS but it didn't engage with regard to the risk). Agule claims that one benefit of a capacitarian account of responsibility is that one is not forced to adopt a counterfactual approach to handle negligence cases. Contra Moore, Agule claims our approach rests responsibility on a counterfactual understanding of something that should have happened, instead of finding a state of affairs within the agent that actually did happen to justify responsibility. In such cases it would seem guidance control is lacking, at least with regard to the possible risk of harm; and yet we want to hold someone responsible anyway.

As our discussion of Moore above may suggest, we think Agule's worries about counterfactual language are overblown. We do not see counterfactual language as incompatible with, or even deeply different from, a sensible understanding of mechanisms and what they can do, and hence feel no need to purify our reasons-responsive account of any and all counterfactual language. Even so, Agule offers us a way to strengthen our theory's handling of negligence cases by specifying what it means to possess a synchronic MWS at the time negligent harm is caused. He indicates we ought to hold a person responsible if (1) they have an MWS, (2) they act or omit to act, and (3) their MWS plays a role in their action or omission (Agule forthcoming, 9). Agule notes that this means the sort of awareness necessary to purposely or knowingly commit crime seems to have to do with the allocation of attentional resources (Agule citing Murray, 10). In negligence cases these resources are not illuminating information about the risk, but this does not mean executive functions are "offline" with regard to an act. If a person is not paying attention to such risks, synchronic inhibition based upon assessment of that risk is impossible, even though the person is, in an important sense, in control of the action that caused legal harm. Obviously, it is possible to perform risky behaviors without being consciously aware of the risk at the time the act is committed.

Agule says that on our executive theory, agents in cases like Montgomery's may be understood as responsible for their negligent wrongdoing where their executive functions are "engaged with" the risk in some—possibly unconscious—sense, but attentional focus is not on the risk (Agule, 11). We might see Montgomery's MWS "engaged with" the risk by noting the quick task-switching from her conversation to her driving and back again, indicating that her cognitive system recognizes the risk of becoming completely absorbed by her conversation. Agule notes we ought not to privilege attention as a part of an MWS necessary to responsibility, and we agree, to an extent. Attention to certain reasons or facts, again, is important to synchronic inhibition based upon those reasons or facts. And this is important to the highest levels of responsibility. But we still do have guidance control in cases of negligence, and importantly, a functioning MWS gives us diachronic control over what we pay attention to and how many and what types of risks we are likely to take. The claim that "the more executive functioning there is, the more agential engagement there is, and thus the more culpable wrongdoing" seems right (Agule forthcoming, 12). Lesser levels of responsibility seem appropriate with lesser levels of engagement between MWS and beliefs relevant to the assessment of risk.

In some cases, however, it seems that “the agent never realizes their behavior is risky” (13). Even so, if the agent possesses a sufficient minimal set of executive functions, “the agent’s executive functioning takes up the relevant evidence, and the agent’s executive system guides the agent’s behavior” (13). To handle these cases, Agule further amends our theory further to require: (1) an MWS, (2) an action or omission, and (3) that the MWS played a role in the action or omission by being engaged with background beliefs relevant to the risk, and perceptions of the features of the situation relevant to the risk. This requires no actual engagement with a perception of risk, unconscious—as exhibited by Montgomery’s quick task-switching—or not. Instead, Montgomery need only have background beliefs such as “It is dangerous to quickly cross two lanes of traffic without looking,” and “I am crossing two lanes of traffic quickly without looking now,” that could have given rise to a perception of risk. This move seems to take us full circle back to a Stark-like position: We can hold someone responsible for negligence when they have background beliefs that could have led to an understanding of the risk, even where the belief regarding the risk is not actually formed, and so long as they have the requisite level of doxastic self-control (Stark 2016, pp. 247–249). In this way Agule helpfully spells out in clear terms what it means for an agent to possess an MWS at the time of a negligent act, and why it matters to responsibility. However, on Agule’s analysis, persons without synchronic possession of an MWS—but only diachronic—would not be responsible. Again, we certainly agree that these are the “easier” cases of negligence, although we argue that in some rare cases, diachronic possession of an MWS may be enough to ground responsibility, because one’s MWS had engaged with the risk and diachronic possession of an MWS means one could have exercised self-control over time to avoid the risk.

2 Does Our Theory Have Novelty/Normative Import?

Our commenters, especially Morse and Patterson, rightly want more than just a good fit between our theory of MWS and legal categories of responsibility. They want our theory to do explanatory work and offer the possibility of a normative effect. Instead, they claim, our theory is interesting but inert. Patterson notes that, at its core, negligence law is about holding persons responsible for violating a standard of care where that violation results in legal harm. We are responsible for such violations even where they involve a failure to pay attention or take note of a risk. Why appeal to an MWS at all, says Patterson? Why not just allude to diachronic possession of capacity responsibility in Hart’s sense? Similarly, Morse is extremely skeptical that facts about executive functions and the brain processes that realize them can offer anything “new” in the way of normative implications beyond what behavioral analysis can provide.

We think our discussion of negligence revealed some interesting things, especially that the justification for holding persons responsible for lapses or inattention may be somewhat different than for voluntary intoxication; we note that Hart’s analysis does not shed light on these differences. But we think our best chance of

showing the potential normative impact of our theory is to focus on our views of juvenile culpability, and we will focus on that topic in this part of our response.

Morse finds little that is novel or valuable in our discussion of juveniles. The immaturity excuse and juvenile courts are longstanding features of the criminal law, and in any case neuroscience is not the place to look, he argues—we already know from the way young people behave that they are immature agents sometimes entitled to leniency, so what does a more granular knowledge of the maturing brain add to that picture?

We think Morse may be missing the primary thrust of our discussion of juvenile justice in RB. The main point of applying our theory of responsibility to juvenile agents is not to make a bold new set of claims that *need* supporting evidence, but for emerging facts about juvenile cognition to serve as supporting evidence for our core theory. Morse is quite right that excuses based on immaturity are a long-established fixture of the criminal law, and that developmental psychologists, legal scholars, and philosophers have already staked out the territory of why juveniles often receive (and deserve) a break. What we hope to achieve, however, is to get beyond generalities about how juvenile agency compares to its adult counterpart, and even to push against the juvenile–adult binary that may obscure more than it reveals. This is where we think neuroscientific findings can have important normative implications about the responsibility of young people.

RB's discussion of juvenile justice aims to make finer-grained distinctions among agents and classes of agents through the identification of the specific brain mechanisms that enable responsible agency. Understanding the complexities of how these mechanisms develop and mature, and how they manifest in behavior, is not something that we would expect to completely revolutionize the juvenile justice system. A better understanding of brain maturation does not “change everything,” and we would certainly want to resist the “lobbyist neuroscience” that Francis Shen warns about (2013, 999). But the choice is not between “changing everything” and changing *nothing*. Our relatively modest hope is simply that our theory enables a more fine-grained and accurate assessment of culpability among juvenile offenders; and our contention is that this understanding calls for them to be treated, in general, with greater leniency.

In RB we call for some specific changes to the current way in which juvenile offenders are handled. Currently, persons who commit a crime under the age of 18 are handled by the juvenile court system unless they get transferred to adult court. Jurisdictions differ, but in the US cases most likely to get transferred are the most serious ones—that is, a murder or armed robbery is much more likely to transfer to adult court than theft. We disagree with transfer based on category of crime—as many do—but also recommend a scalar approach to responsibility and a graded structure of legal categories to give a finer-grained and more accurate assessment of how juvenile offenders develop toward full, adult-level responsibility. We do so on the basis of emerging science indicating the rate at which different executive functions become fully developed.

On our view, offenders under the age of 12 would be under an irrebuttable presumption of incapacity on grounds of immaturity; those from 13 to 17, under an irrebuttable presumption of *diminished* capacity on grounds of immaturity; and

those from ages 18 to 21, under a rebuttable presumption of diminished capacity on grounds of immaturity. From the age of 22 onward offenders would be under a presumption of adult-level capacity that cannot be rebutted on grounds of immaturity, though not one that is irrebuttable *simpliciter*. To complement this reform of the guilt phase, we should also rethink the dominant approaches to punishing juveniles at the sentencing phase—when interventions are appropriate, they ought to reflect the reduced culpability of juvenile offenders and promote rather than (further) arrest the development of executive capacities. Enacting these proposals would, obviously, entail a comprehensive revolution of the juvenile justice system—not changing *everything*, again, but far from changing nothing.

Morse (this volume) disagrees with our proposal, but his main complaint has less to do with its substance or justification and more to do with the fact that our proposal is, in his telling, “effectively based entirely on behavioral analysis despite the apparently obligatory invocations of brain maturation”. His question of what all this “brain talk” comes to is well taken—it calls to mind a past remark of Morse’s on the same topic of juvenile justice: “If the behavioral differences between adolescents and adults were slight, it would not matter if their brains were quite different. Similarly, if the behavioral differences were sufficient for moral and constitutional differential treatment, then it would not matter if the brains were essentially indistinguishable” (Morse and Roskies 2013, xxii).

We grant that our understanding of executive function in juveniles is informed by behavioral analysis. Because the criminal law speaks in the language of folk psychology, neuroscientific results must be linked to behaviors and behavioral capacities if they are to illuminate aspects of responsible agency in juveniles. Surely, the mere inclusion of behavioral analysis with neuroscience in an approach to juvenile responsibility—along with attention to the broadening life experiences of adolescents, the increasing range of reasons to which maturing juveniles can become sensitive, and more—does not render the neuroscientific findings empty or otiose.

Morse’s worries are legitimate, and we appreciate being pressed from his direction on the novelty and normative import of our theory. But we find this line of critique to be hampered by persistent false dichotomies, including the suggestion that an approach using both neuroscientific and behavioral data is *simply* “behavioral analysis,” as if one could either have a neuroscientifically informed approach to juvenile justice or a behaviorally informed one, but not one that was both.

This last point raises the question that is at the heart of Morse’s and Patterson’s comments: If something like our tiered structure of juvenile responsibility is a good idea, isn’t that already sufficiently established by the psychological and behavioral evidence on offer? What exactly does the neuroscience *add*? Without overstating the case or engaging in “neurohype,” let us sketch five or six possible, provisional answers.

First, neuroscience complicates and enriches the emerging view of when, on average, young people reach adult-level normative competence. If our moral and legal practices deem it important to discount culpability in light of immaturity, then the details of how maturation occurs can matter a great deal, and neuroscience furnishes a key source (though not the only source) of data on that score. Attending to the neuroscientific details can help cure legal scholars and practitioners of the

simplistic, potentially misleading idea of a unified, progressive trajectory of “brain development,” because, as Somerville (2016) argues, maturation occurs at several levels simultaneously. One can speak of maturity with respect to structure in isolation from maturity with respect to, for instance, function or intrinsic connectivity. The “maturational asymptotes” for some of these levels appear to converge; others appear to diverge; furthermore, each of these levels itself comprises multiple indices of maturation. Even focusing on functional maturity alone, “there is no single progression” to be found. Rather, “neural activity intensifies and reduces, varies quantitatively and qualitatively, in linear and nonlinear ways that are both linked to—and independent of—behavioral differences across development” (Somerville 2016, p. 1166).

Having a better sense of these complexities would mean a better sense of what certain neuroimaging evidence might—or might not—suggest about a given juvenile’s psychological and behavioral capacities. Evidence of a functional or structural abnormality in a juvenile offender’s brain might be highly relevant to assessing his normative competence relative to his chronological age; it seems equally likely, however, that a single index of neurobiological maturity would be (reasonably!) judged insufficient, if not irrelevant, to questions of normative competence. If one is worried that neuroscientific results will be applied in reckless or overzealous ways to matters of law and policy, then more engagement with these results in all their complexity is the best way forward. The cure for “brain overclaim syndrome,” in Morse’s (2006) memorable turn of phrase, is not less neuroscience but more of it.

A more accurate, nuanced, and complete picture of the “multiple maturities” (Somerville 2016) that underwrite normative competence should also help push back against the strict binary between juveniles and adults that continues to hold sway in legal and scientific contexts. This binary has some obvious usefulness, but it also has a circular, self-justifying logic: Scientific studies often use 18 as the “cut-off” age between adolescence and adulthood, partly because that age is so frequently the relevant legal division between juveniles and adults; the findings from such studies, even if they have novel implications in some respects, implicitly reinforce the legitimacy of that legal cutoff, which in turn justifies the cutoff being used in further scientific studies, and so on. Breaking this unhelpful cycle will take efforts within the legal world and the scientific one, as well as thoughtful communication between the two. The age-of-18 cutoff does not necessarily have to be thrown out as a useless and arbitrary boundary in either field, but both would be better served by moving toward a model comprising *several* “cutoffs” between phases of adolescence and young adulthood, allowing both courts and scientists to make finer distinctions among a population that is undergoing a dynamic and complex process of neural and psychological development that tends not to “level off” until well into the twenties.

Along these lines, another promising contribution from neuroscience comes from several ongoing large, multimodal brain imaging projects, which hold the promise of establishing standards, backed by huge datasets, for normative brain development. These projects, such as the Lifespan Human Connectome Project in Development (Harms et al. 2018, enrolling over 1,300 subjects aged 5–21) and the Adolescent Brain Cognitive Development study (Casey et al. 2018, including over 10,000 adolescents in a longitudinal study), are still in early stages, and their leaders are careful

to note the limits of their findings. Still, the prospect of establishing multimodal neurocognitive “growth curves”—not unlike the familiar ones used by pediatricians to put your child’s height and weight in context—has obvious relevance for juvenile justice. As neuroimaging evidence increasingly finds its way into courtrooms in ways that may either illuminate or mislead (Farahany 2015; Kuersten 2016), advocates, judges, and juries may be prompted to consider what healthy brain development looks like in a typical 9-year-old, 15-year-old, or 20-year-old. Only high-powered studies like the ones above offer the hope of an answer.

Another application of neuroscience in this context would be determining relative similarities between age-range groups. Offenders in their late teens and early twenties are a group of particular focus in juvenile justice, in part because they cluster around the traditional cutoff age of adult criminal responsibility. Even if we are speaking quite generally or in terms of “average” members of these groups, it would be worth knowing whether, with regard to the structures and processes that subserve executive function, the average 20-year-old’s brain is more similar to that of an average 15-year-old or an average 25-year-old, despite being chronologically equidistant from both.

Establishing a general picture of what “average” or “typical” brains look like at various stages of development is, as the foregoing examples indicate, highly relevant to questions of juvenile justice. But generalizations only go so far; each young person’s brain follows a unique “journey” of development in a unique social and experiential context. This is all the more reason, however, to sweat the neuroscientific details. In this spirit, another of our commenters, Federica Coppola (2018), articulates some of the ways that neuroscience could inform practices of sentencing and treatment, by using “neuroprediction” to model and assess young offenders’ risk of future criminal behavior and tailor interventions to fit a particular young person’s situation. Naturally this proposal comes with a host of legal and moral concerns, yet Coppola offers a compelling account of how those concerns can be addressed in the pursuit of an approach to sentencing that is both less retributive and more finely tuned to the needs and developmental trajectories of individual offenders. When used carefully—and always, of course, in conjunction with psychological and behavioral data—developmental neuroscience can help us discern, for each individual juvenile offender, what sorts of interventions are developmentally appropriate and well-timed to foster rehabilitation and social reintegration (Coppola 2018).

Some of these possibilities may be a long way off yet—some require a more mature neuroscience, and a far more robust and well-settled body of findings. But these are chiefly practical problems to be settled in the lab and the scanner; they are not reflective of in-principle limits on what neuroscience might reveal, suggest, or imply regarding juvenile responsibility. The theory of juvenile culpability promoted in RB provides a framework that we think reflects the current state of things regarding developmental neuroscience. However, our theory is well-situated to accommodate emerging neuroscience, to the point where our presumptions of incapacity and capacity may shift in light of new evidence.

Although we have primarily drawn on Morse’s remarks in this section, we would argue that many of these points apply to Patterson’s (2021) criticisms of inertness as well. Patterson alleges that our account offers, at best, only “an elaborate explication

of what we already know” (9). This remark echoes Justice Kennedy’s opinion in *Roper*, wherein he noted that juveniles’ reduced culpability is part of the common wisdom that “any parent knows.” But when it comes to something like the delicate and complex process by which a non-responsible child develops into a fully-responsible adult—and the way a massive system of laws and punishments handles that process—the details of “what we already know” matter a great deal and explicating them can have significant normative impact on judgments of culpability and practices of punishment. Understanding how the distinct-but-never-entirely-separate executive functions come online, develop, and gradually expand the reasons-responsiveness of juveniles led us to recommend important revisions to the way juvenile offenders are handled, so we do not think our theory is simply a matter of filling in the specifics of a general theory already set down by Hart and others. But even if this were the case, it would hardly make those specifics inert: for example, new specifics may mean new measures of neurological and psychological maturity are taken as evidence of responsible agency. This point generalizes to cover not only our discussion of juvenile justice, but cases involving legal insanity, psychopathy, negligence, and more. The value of our project lies in connecting, e.g., Hart’s broad (and broadly correct!) views about capacity and responsibility to a detailed neuroscientific account of how executive functions work in different kinds of brains; the normative implications of that *connection*, if not of the neuroscientific facts themselves, are potentially profound.

References

- Agule, Craig. 2021. “Minding Negligence.” *Criminal Law and Philosophy*.
- Baron, Marcia. 2020. “Negligence, Mens Rea, and What We Want the Element of Mens Rea to Provide.” *Criminal Law and Philosophy* 14 (1): 69–89.
- Casey, B.J., Tariq Cannonier, May I. Conley, Alexandra O. Cohen, Deanna M. Barch, Mary M. Heitzeg, Mary E. Soules, Theresa Teslovich, Danielle V. Dellarco, and Hugh Garavan. 2018. “The Adolescent Brain Cognitive Development (ABCD) Study: Imaging Acquisition across 21 Sites.” *Developmental cognitive neuroscience* 32: 43–54.
- Coppola, Federica. 2018. “Mapping the Brain to Predict Antisocial Behaviour: New Frontiers in Neurocriminology, ‘New’ Challenges for Criminal Justice.” *UCL Journal of Law and Jurisprudence Special Issue* 1 (1): 103–126.
- Farahany, Nita A. 2015. “Neuroscience and Behavioral Genetics in US Criminal Law: An Empirical Analysis.” *Journal of Law and the Biosciences* 2 (3): 485–509.
- Fischer, John Martin, and Mark Ravizza. 1998. *Responsibility and Control: A Theory of Moral Responsibility*. Cambridge: Cambridge University Press.
- Fischer, John Martin, and Neal A. Tognazzini. 2009. “The Truth about Tracing.” *Noûs* 43 (3): 531–556.
- Frankfurt, Harry G. 1969. “Alternate Possibilities and Moral Responsibility.” *The Journal of Philosophy* 66 (23): 829–839.
- Garvey, Stephen P. 2006. “What’s Wrong with Involuntary Manslaughter?” *Texas Law Review* 85 (2): 333–384.
- Harms, Michael P., Leah H. Somerville, Beau M. Ances, Jesper Andersson, Deanna M. Barch, Matteo Bastiani, Susan Y. Bookheimer, Timothy B. Brown, Randy L. Buckner, and Gregory C. Burgess. 2018. “Extending the Human Connectome Project across Ages: Imaging Protocols for the Lifespan Development and Aging Projects.” *Neuroimage* 183: 972–984.
- Hart, H.L.A. 1968. *Punishment and Responsibility: Essays in the Philosophy of Law*. New York: Oxford University Press.

- Hirstein, William, Katrina L. Sifferd, and Tyler K. Fagan. 2018. *Responsible Brains: Neuroscience, Law, and Human Culpability*. Cambridge, MA: MIT Press.
- Husak, Douglas. 2016. *Ignorance of Law: A Philosophical Inquiry*. New York: Oxford University Press.
- Husak, Douglas.. 2021. "The Objective(s) of Responsible Brains." *Criminal Law and Philosophy*.
- King, Matt. 2009. "The Trouble with Negligence." *Social Theory and Practice* 35 (4): 577–595.
- Kuersten, Andreas. 2016. "When a Picture Is Not Worth a Thousand Words." *Geo. Wash. L. Rev. Arguendo* 84: 178.
- Levy, Neil. 2014. *Consciousness and Moral Responsibility*. Oxford: Oxford University Press.
- McGeer, Victoria. 2015. "Building a Better Theory of Responsibility." *Philosophical Studies* 172 (10): 2635–2649.
- McGeer, Victoria, and Philip Pettit. 2013. "The Hard Problem of Responsibility." In *Oxford Studies in Agency and Responsibility*, edited by David Shoemaker. Oxford: Oxford University Press.
- Moore, Michael S. 2021. "Relating Neuroscience to Responsibility: Comments on Hirstein, Sifferd, and Fagan's *Responsible Brains*." *Criminal Law and Philosophy*.
- Moore, Michael S., and Heidi M. Hurd. 2011. "Punishing the Awkward, the Stupid, the Weak, and the Selfish: The Culpability of Negligence." *Criminal Law and Philosophy* 5 (2): 147–198.
- Morse, Stephen J. 2006. "Brain Overclaim Syndrome and Criminal Responsibility: A Diagnostic Note." *Ohio St. J. Crim. L.* 3: 397–543.
- Morse, Stephen J. 2021. "Is Executive Function the Universal Acid?" *Criminal Law and Philosophy*.
- Morse, Stephen J., and Adina L. Roskies. 2013. *A Primer on Criminal Law and Neuroscience*. Oxford: Oxford University Press.
- Patterson, Dennis. 2021. "Inert." *Criminal Law and Philosophy*.
- Roskies, Adina L. 2018. "Decision-Making and Self-Governing Systems." *Neuroethics* 11 (3): 245–257.
- Sarch, Alexander. 2018. "Review of Findlay Stark, *Culpable Carelessness: Recklessness and Negligence in the Criminal Law*." *Criminal Law and Philosophy* 12 (4): 725–730.
- Shen, Francis X. 2013. "Legislating Neuroscience: The Case of Juvenile Justice." *Loyola of Los Angeles Law Review* 46 (3): 985–1018.
- Sifferd, Katrina L. 2018. "Ought Ignorance of Criminal Law (and Its Morality) Excuse? An essay review of Husak, D. *Ignorance of Law: A Philosophical Inquiry*." *Jurisprudence* 9 (1): 186–191.
- Sifferd, Katrina L. Forthcoming. "How Is Criminal Punishment Forward-Looking?" *The Monist*.
- Somerville, Leah H. 2016. "Searching for Signatures of Brain Maturity: What Are We Searching for?" *Neuron* 92 (6): 1164–1167.
- Stark, Findlay. 2016. *Culpable Carelessness: Recklessness and Negligence in the Criminal Law*. New York: Cambridge University Press.
- Vargas, Manuel. 2005. "The Trouble with Tracing." *Midwest Studies in Philosophy* 29 (1): 269–291.
- Vargas, Manuel. 2013. *Building Better Beings: A Theory of Moral Responsibility*. Oxford: Oxford University Press.
- Vargas, Manuel. 2015. "Precis of Building Better Beings: A Theory of Moral Responsibility." *Philosophical Studies* 172: 2621–2623.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.