

# Tracing the Epistemic Condition

*Matt King*

## 1. Anatomy of Tracing for Action

There are cases in which an agent appears responsible for their action despite seemingly lacking the requisite control. A canonical example is drunk driving. Suppose Topsy Tina has had too much to drink, yet drives home. On the way, she hits a family sedan, injuring those inside. By hypothesis, Tina is so intoxicated as to lack the control over her actions ordinarily required to be responsible. Nevertheless, Tina is surely blameworthy for having injured the family.

To reconcile these two claims, some theorists have turned to the notion of tracing.<sup>1</sup> In short, we can trace back to an earlier action of Tina's that has the following two properties: (1) she exercised sufficient control over that action to be responsible for it, and (2) it explains her lack of control for the subsequent action. In Tina's case, she can be said to have controlled her getting drunk, which explains why she was uncontrolled at the time of the crash. These facts allow us to "trace" her responsibility (and, indeed, blameworthiness) for the crash back to her responsibility for getting drunk.

Explanations of cases like Tina's via tracing amount to exceptional explanations of responsibility. Tracing isn't needed to explain responsibility in cases of ordinary action, like making waffles for breakfast or when Tina exercises sufficient control over parking her car at the bar. Instead, tracing is a special explanatory mechanism to be used in cases of apparent responsibility despite the lack of control, and it applies wherever those two conditions obtain.

I have argued elsewhere that accounts of the control required for responsibility can explain cases like Tina's without any appeal to tracing (2014). In brief, the strategy is to show that we can account for the relevant sorts of cases by appeal to ordinary explanatory mechanisms. In particular, I argued there that we could explain responsibility in drunk-driving cases on either a recklessness model or a negligence model,

<sup>1</sup> For positive defenses of the approach, see Fischer and Ravizza (1998) and Mele (1995), though each targets phenomena more diverse than just drunk-driving cases. For good literature reviews as well as discussion of various complications, see McKenna (2008) and Vargas (2005).

but that on neither model was tracing a necessary component. This chapter considers the prospects for that general line of argument with respect to tracing applied to the epistemic condition on responsibility. I intend to remain neutral on how best to understand such a condition and how independent from control it is. Instead, I propose to consider epistemic dimensions of responsibility, where an agent's beliefs (or similar cognitive states) play a particular role in explaining why they are or are not responsible, irrespective of how best to incorporate this dimension into our overall conditions on responsibility.<sup>2</sup>

I believe that tracing remains a dispensable feature of our accounts of responsibility.<sup>3</sup> In the following sections, I draw out how I understand tracing to apply to the epistemic dimension and argue we need no special explanatory mechanism like tracing to explain responsibility and blameworthiness.<sup>4</sup> Additionally, I argue that, in some instances, tracing as an explanation is in fact inadequate.

## 2. Ignorance and Plans

Ignorance and mistakes often excuse. If Benson doesn't know that the button he pushes to start his coffee maker in the morning is hooked up to a bomb, then he isn't responsible for blowing up the bakery. We can explain the relevance of his ignorance in various ways. For the purposes of this chapter, we can adopt a simple working view that Benson's lack of information shows that he didn't blow up the bakery intentionally, on a sufficiently broad reading of that term. He didn't do it on purpose or knowingly.

But what if Benson had known earlier about the button. In fact, he himself set it up last night so that it would detonate the bomb when he made his morning coffee. However, he also took a pill shortly thereafter, which eliminated those memories. He subsequently forgot about his plan to blow up the bakery, about planting the explosive, wiring the coffee maker, etc. He intended to erase the relevant epistemic states, such that, as a result, when he presses the button in the morning, he no longer knows that he is not just making coffee, but also setting off a bomb. Thus, Benson intentionally brings it about that he unknowingly blows up the bakery.<sup>5</sup>

<sup>2</sup> Thus, I intend to elide questions of whether there is an independent epistemic condition on responsibility. See Björnsson (2017) and Graham (2017). Relatedly, I intend to elide questions of whether moral ignorance (i.e., ignorance or mistakes explicitly about the moral character of one's action) should be treated differently from non-normative ignorance or mistakes. See Harman (2017) and Alvarez and Littlejohn (2017).

<sup>3</sup> Contra Fischer and Tognazzini (2009), who claim that tracing is an indispensable feature of theories of responsibility.

<sup>4</sup> Clarke (2017) also defends the claim that we can be directly responsible for at least some actions out of ignorance, and does so via appeal to ordinary explanatory mechanisms. However, it is less clear whether his account can forgo reliance on tracing altogether.

<sup>5</sup> To keep the rationality of his actions intact, we can assume it leaves his motivating desires intact, so that he is still pleased the bakery exploded, though he is unaware of his role in its destruction.

I take it that Benson is still blameworthy for the bomb's destruction. But he is ignorant of the fact that that's what he's doing when he turns on the coffee maker. Therefore, if Benson is blameworthy for blowing up the bakery though ignorant, it must be that his ignorance does not excuse.

One way we might account for this disparity is via tracing. Since Benson is responsible for his own ignorance, we can connect his responsibility for the explosion back to his responsibility for that earlier episode. Since Benson, as it were, intentionally made himself ignorant, he is responsible for doing so. Since it's his own fault that he doesn't know he's exploding a bomb in the morning, he can't make use of the excuse of ignorance. Instead, we can trace his responsibility for the explosion to his responsibility for making himself ignorant. We can turn to tracing to account for blameworthiness in the face of apparent ignorance.

But *must* we appeal to tracing? I don't think so. To see why, consider Benson's activity leading up to the explosion. He forms a plan: to set up a process by which a bomb will explode. And this is the plan that he enacts. He then performs a number of actions in service of that plan. He wires elements together, packages an explosive, and places the bomb somewhere in the bakery. When he switches the coffee maker on, he detonates the bomb. I believe that Benson detonates the bomb knowingly, despite the fact that, at the time of the explosion, he is unaware that he is causing an explosion.

This might seem paradoxical. How can Benson knowingly blow up the bakery when he doesn't know at the time that he is detonating anything? But there is really no mystery here. Agents are capable of quite sophisticated planning. Their abilities include not only to achieve goals but to achieve them in particular ways. In this case, Benson plans not only to blow up the bakery, but to do so in such a way that he isn't aware at the time that he is doing it. Such is the power of agents.<sup>6</sup>

Indeed, technology facilitates such actions all the time. I can set up my email to send out a message at a particular time. I haven't emailed anyone until the message is sent, and it is possible that when the message is sent, I will be unaware that I am sending it. Perhaps I forgot about it altogether. Nonetheless, it would abuse the language to say I unintentionally send it out, or emailed by mistake or under ignorance. In fact, my tendency to forget about having to send the email is precisely why I scheduled it to happen in the first place. I set a plan in motion so as to, in some sense, insure my later actions against my intervening psychology.

The same can be said of Benson. Perhaps he feared his resolve would weaken overnight, and come morning he'd abandon his plan. So he took steps to protect his plan from failing, just as one might take steps to guard against interference from an outside threat. More mundane examples abound. I routinely set my dishwasher to start on an hour delay, at least in part because of my tendency to later forget to start it

<sup>6</sup> For a related discussion, see Gideon Yaffe's *Attempts* (2011, ch. 6), wherein he argues that one can intentionally do the unintentional as a way of showing how attempt liability for crimes requiring a *mens rea* of negligence is possible.

otherwise. I run it intentionally each time, regardless of my occurrent mental states when its cycle begins.

One might object that my account so far is simply a redescription of tracing. There is something I am responsible for (setting the dishwasher to run) that we trace back to in order to establish responsibility for some later thing (running the dishwasher). But it is not enough for the tracing theorist to show that an agent can be responsible for something in the past that (partially) explains their responsibility at present. All action is arguably like this. For most actions, there is some other action, which occurred prior, such that we could include it in part of the explanation for the agent's responsibility for the later action. Suppose I break a window by hitting it with a hammer. My breaking of the window is at least partially due to my swinging the hammer at it (just as my swinging the hammer forward is partially due to my swinging it back). But it would be otiose to insist that my responsibility for breaking the window is due to my responsibility for swinging the hammer.<sup>7</sup>

This is not to say we cannot explain the responsibility of the one by tracing back to responsibility for the other.<sup>8</sup> But my aim here isn't to show that tracing is explanatorily *inadequate*, only that it is explanatorily *unnecessary*.<sup>9</sup> That is, there is no special explanatory mechanism required to go by the name "tracing." Part of the reason why is that we don't require it to explain ordinary action, like breaking windows or opening doors (the responsibility for which could similarly be explained by responsibility for turning the knob and pulling).

We can extend the account covering ordinary actions to cases often cited as requiring tracing, however. This extension is possible because of how our mental states help organize our actions over time. Consider temporally extended actions, like taking a walk or baking a cake.<sup>10</sup> We may interrupt such activities with other actions without ceasing the original activity altogether. If one stops to smell the roses, one has

<sup>7</sup> Some action theorists may insist that complex actions, like running the dishwasher, are to be decomposed into separate constituent actions. Others might insist that we distinguish between actions and outcomes, where the latter are products of the former. Either claim might seemingly require tracing in order to connect the constituent action-components or outcomes to their actions. While I do not dispute that we can carve up actions in this way, as well as distinguishing them from outcomes, I see no principled reason that requires doing so. (And certainly no principled reason to require it in advance of the observations of the discussion herein.) Since my aim here is simply to deny the necessity of tracing, I can be taken as equally denying the necessity of decomposing complex actions as well as forcing a distinction between actions and outcomes.

<sup>8</sup> Though commitment to such a transfer principle would arguably open the door to skeptical worries about the very possibility of responsibility (cf. Strawson 1994). I am here assuming skepticism about moral responsibility, whether generally or just for unwitting acts, is false. See Levy (2017) for discussion of the latter. Skeptics of either kind are invited to treat my argument here as conditional: if we can be responsible for unwitting acts, we can explain such facts without an appeal to tracing.

<sup>9</sup> At least in the relevant sorts of cases discussed here. I leave unresolved whether there are further contexts in which tracing might yet be required.

<sup>10</sup> Of course, few actions aren't somewhat temporally extended, as the hammer-swinging illustrates.

not thereby stopped taking a walk.<sup>11</sup> Or one may be in the middle of baking a cake and have to run out to the grocery store for a missing ingredient. It is accurate when seeking assistance in the store to say, “Well, I’m baking a cake and need vanilla.” You are still baking a cake despite the fact that you are at the grocery store.<sup>12</sup> Alternatively, one may stop one’s walk by stopping to smell the roses. Or one might cease making a cake after running into a friend and deciding to go get coffee together. What makes the difference in each case are the agent’s organizing mental states, the persisting intentions or beliefs about one’s own action.<sup>13</sup> There is no need to carve up the components of each complex action into its constituents and trace responsibility throughout. Rather, the whole can remain, and be fully intentional, even as the actor’s occurrent thinking might fluctuate.

The above is true regardless of whether one intends the action or merely foresees it (or one of its components). Before setting off, Sergio may consider the possibility that he stops to smell the roses during his walk. He proceeds anyhow. When he stops to smell the roses, he may do so absent-mindedly, thinking instead only of the complicated philosophical argument he’s trying to work out in his head. Nonetheless, and though there is no good adverb in English, I think he smells the roses in such a way that would count as reckless were the action unjustified. That is just to say he does so more or less “knowingly.”

I’ll return to recklessness in a moment, but first I want to highlight here that the responsibility ascriptions, while they reference some past fact, do not trace back to a prior responsibility *attribution*. The responsibility for smelling the roses is a contemporaneous matter, even while the fact that it was a “risk” knowingly run is true in part because of past facts. It isn’t that Sergio is responsible for having thought about stopping to smell the roses that explains why he is responsible for so stopping. Rather, the claim is that what explains why his stopping to smell the roses is done knowingly (more or less) is that he had considered that he might do so. Ascriptions of responsibility for that prior consideration are actually mistaken here.<sup>14</sup>

None of this shows that tracing is mistaken. One could instead insist that tracing is simply ubiquitous. Every (or nearly every) instance of responsibility is traceable to

<sup>11</sup> I believe this example, though common enough, is borrowed from Sergio Tenenbaum’s discussion of “gappy” actions (ms).

<sup>12</sup> This is arguably a hallmark of any true description in the progressive. For discussion of similar cases, though not on the subject of tracing, see Falvey (2000).

<sup>13</sup> Cf. Bratman (1987) on the planning capacities of agents. I mean here to rely on the basic observation only, not Bratman’s full theory of intention.

<sup>14</sup> It would be incredibly odd, I think, to suppose that in some piece of deliberation about what to do, we should separately ascribe responsibility for the various aspects of the action or occurrences in thought of possible outcomes, side-effects, or interruptions that might occur. It is of course true that to have considered some aspect of my plan of action it will have to have occurred to me (in some sense of occurred). But this seemingly has nothing to do at all with ascriptions of responsibility. Even if awareness of the possibility of stopping to smell the roses had been implanted in Sergio’s brain by the ubiquitous evil neuroscientist (and nothing more), this would seem sufficient for his having considered it, regardless of the thought’s etiology.

some previous instance of responsibility, at least for any moderately involved action. But notice that if one thinks that tracing is a feature of all (or nearly all) explanations of responsibility, then this actually undercuts a major motivation for it. After all, it seems natural to look to tracing only when we are missing a usual condition on responsibility. Recall that tracing seems needed for cases like Topsy Tina or Benson, cases that are importantly aberrant, missing a condition like control or knowledge. If tracing in fact figures into most (if not all) instances of responsibility, it is unclear why we should need to resort to any special explanatory mechanism whatsoever.<sup>15</sup>

### 3. Foreseeability and Tracing

But what of the aberrant cases? Benson's case may be unusual in its details, but there are clear enough ordinary cases in which agents bring about harm out of supposed culpable ignorance. Consider a popular example from the literature. Suppose a doctor fails to stay up-to-date on the latest medical research relevant to her practice. As a result she causes harm to a patient, harm she would have known how not to cause had she properly informed herself. She could have read the relevant recent medical research, but neglected to.<sup>16</sup>

Should she harm a patient as a result, it may seem obvious that the doctor is blameworthy for harming only because she is culpable for not knowing, and that she is so culpable only because she was blameworthy for failing to inform herself properly. But the question remains whether this difference can only be explained via tracing. If our concern is culpable ignorance, then what is to be shown is that the doctor, though ignorant of the risk of harm, is nonetheless blameworthy for the harm caused *because* she is blameworthy for not knowing about the risk.

On one interpretation of the case, however, the doctor doesn't even look relevantly ignorant. If she actually considered the risk of harm practicing without consulting the research posed, then proceeding without consulting the research is arguably reckless. Suppose that Rose knows that the research is relevant to her future services. If one has reason to think some information one has access to is relevant to the risk one may pose to others, and one does not examine that information, then one proceeds at one's own moral peril. Whether or not accepting that risk is reckless will turn on whether or not running that risk is justified. If Rose has reason to suppose there is some new data on the risk of a minor side-effect of the medication or procedure, perhaps taking on that risk is allowable. If the data concerns the risk of a fatal complication, then that risk looks unjustified.<sup>17</sup>

<sup>15</sup> I return to these considerations in section 6.

<sup>16</sup> This particular example is drawn from Holly Smith's influential "Culpable Ignorance" (1983). Not only is it a good test case from the literature, but it is more closely isomorphic with cases like Tina's, where control is the primary consideration.

<sup>17</sup> Risk varies in response to both the likelihood of harm and its severity, but I take no stand here on a method for evaluating the acceptability of risk.

One who accepts<sup>18</sup> an unjustified risk is reckless, and is responsible for the harm risked when it occurs. As I illustrated above, one can do something knowingly without being aware at the time that one is doing it, just as one can intentionally do something without occurrent awareness. As a result, Rose is blameworthy but not relevantly ignorant. Admittedly, she doesn't know that her treatment will harm her patient. But she also doesn't think it definitely won't. Indeed, she believes it might harm her patient since she's aware of the risk of harm. She recognized that possibility and proceeded without settling the matter or revising her plans.

Understanding the acceptance of risk in this way helps explain why foreseeability should be constrained in the way others have noted.<sup>19</sup> Rose need not appreciate that she will harm this particular patient in just so-and-so a way to be blameworthy for harming this patient in the way that she actually does. All she need foresee is that she might cause harm within the "scope" of the risk.<sup>20</sup> To recognize a risk just is to foresee a range of harms under possible sets of circumstances.<sup>21</sup> So whatever harms Rose foresees are the risks she recognizes, and to the extent that she is alive to various possibilities as she deliberates and plans, these risks will be within the ambit of those plans. Thus, when she acts within the scope of those plans, she risks those harms. When unjustified, those risks make her reckless.

We can explain all these facts, however, without resort to any special explanatory mechanism. What makes her reckless lies partially in the past, true enough, as she makes certain decisions or takes action with certain mental states. But the temporal position of her decision before her causing the harm isn't exceptional. What one does intentionally or recklessly depends on one's epistemic states. And one's epistemic states clearly depend on past epistemic states, actions, and events. For example, her reading the article would explain her knowledge of the increased risk of fatal complication or her reading the chart before her would explain her knowledge of the patient's medical history. But while this is a feature of epistemic status generally, it doesn't necessitate invoking tracing. We don't need our model to posit separate and discrete responsibility attributions.<sup>22</sup> Rather, decisions to proceed recklessly open one up to blameworthiness for those harms whose risk one disregarded. Those relevant harms are done knowingly, in a sense sufficient for ordinary responsibility attributions.

<sup>18</sup> I use "accept" (or "acceptance") minimally; it means only that the agent considered the risk and took it to be insufficient reason to revise her action or plan. She is justified in doing so if it provided insufficient reason to revise; she is unjustified if she should have revised in light of the risk.

<sup>19</sup> E.g. Fischer and Tognazzini (2009); Wieland and Robichaud (2017).

<sup>20</sup> Similarly, the doctor need not foresee that they will harm a *particular* patient to be blameworthy for harming that very patient. Understanding the scope of risk in this context is introduced in Smith (1983) (though I disagree somewhat with how she interprets that scope).

<sup>21</sup> Though that range of harms might be limited to a single kind and the set of circumstances could be rather circumscribed.

<sup>22</sup> Indeed, doing so would actually seem mistaken (see earlier).

Why think that Rose causes harm knowingly? For the reasons so far considered, she cannot dismiss the harm as being outside the purview of her plan.<sup>23</sup> It would be infelicitous of her to claim she didn't mean to harm her patients. Of course, she didn't explicitly mean to harm them, but she certainly didn't mean to *not* harm them, since she took no precautions to avoid the risk of harm. Moreover, she's in precisely the same *epistemic* relation to the harm caused as a malevolent doctor who wants to harm their patient, but who knows only that they risk such harm in continuing with the treatment. The difference between the two doctors concerns their attitudes toward the harms involved, partially reflected in the structure of their plans. The latter is positively seeking the risk, and so not reading the research so as to make the harm more likely; whereas Rose is only indifferent to the risk, inasmuch as she doesn't take it to be significant enough to revise her plan. Still, the harm risked is within the scope of Rose's plan in the same way as in the malevolent doctor's. Each has the same evidence and knowledge with respect to it. Just as Benson constructed a set of actions to bring about a particular result (the destruction of the bakery), Rose constructed a set of actions to bring about a particular result (treating her patient) while knowingly risking another (harming her patient).

#### 4. Negligence and Tracing

So far, the discussion has focused on cases in which agents more or less decide not to improve their epistemic position. (In some cases, like Benson's bombing, the decision is to positively degrade it.) I have argued that responsibility in such cases is suitably explained through ordinary mechanisms of intentional and reckless action. I have also claimed that under such interpretations, cases of supposed culpable ignorance turn out to not really involve ignorance of the relevant sort.

But many cases won't involve a prior consideration as to the risk of harm. Recall that our target scenario at present involves a doctor who fails to inform herself of recent medical research, research which, had she read it, would have led her to avoid causing harm to a patient (which she in fact causes). While she may have decided not to consult the research and foreseen the risk of harm as a result, it is more plausible to suppose that she just doesn't think about the risk of harm at all; indeed, she might not even decide to ignore the research. Instead, she might simply fail to think of informing herself, though she should. She may know that a particular research journal has arrived in the office, but, though she should consult it at some point this week, she set it aside when it arrived and has failed to consider it again. Such a doctor is arguably negligent, and negligent agents are responsible for the harms their negligence

<sup>23</sup> To reiterate, "plans" here designate only thin, temporally extended intentional states that organize our actions over time. They needn't be robust and complex affairs, as is the more colloquial notion.

brings about.<sup>24</sup> The new question, then, is whether negligence requires tracing. I do not think so. In fact, I think explaining negligence via tracing is problematic.

To see why, we need to recall how tracing is meant to work. We take an instance where an element required for responsibility is absent, and trace backward to a previous instance where that element is satisfied, and which is related in the right way to the present instance of interest. But such a model cannot work for negligence. If the requisite element is that the risk is at least foreseen, or some other positive epistemic state, then, while we could trace backward to look for it, we will never find it.

Suppose a different doctor, Jo, didn't even think to read the recent medical research. Of course she knows that she should in a dispositional sense, but she doesn't opt to ignore the opportunity to better inform herself. Rather, staying up-to-date is just a standing duty of physicians in her position. Jo is negligent in causing the harm to the patient so long as she should have realized the unjustified risk she was posing her patient.<sup>25</sup> Negligence does not require Jo to actually recognize the risk of harm, only for it to be the case that she should have recognized it or taken the requisite steps to prevent or mitigate that harm. Naturally, those steps in this instance would have been to better inform herself on the latest research and pursue an alternative treatment. By hypothesis, Jo doesn't even think to consider whether she should consult the latest research.

But that means we've traced back from her failure to avoid risking harm to the patient to a failure to consult the latest research. But, epistemically, these two failures are of a piece. In both instances, Jo fails to recognize a risk she ought to recognize without realizing she's done so. Negligence here just traces back to negligence. If that negligence were sufficient to explain responsibility, then we wouldn't need to trace back, because she is already negligent in treating the patient. If negligence is insufficient, then tracing back to a further instance of negligence cannot help ground responsibility either.

Thus, negligence poses a dilemma for tracing. Either it can explain responsibility (in our case) at the time of treatment and so without appeal to tracing, or if it can't, tracing won't be up to the task, for the most immediate instances to trace back to will be further instances of negligence.<sup>26</sup>

Some might insist that tracing preserves a difference between Jo's case and a similar doctor who also causes harm, but who wasn't culpable for their ignorance. Suppose Joaquín regularly keeps up with the latest research, but the particular journal with the pertinent information was delayed in the mail (and so a day late). Supposing he treats a patient under identical circumstances as Jo, he does not seem

<sup>24</sup> I have elsewhere defended a skeptical line with respect to negligence (2009). But I set such skeptical worries aside here for argument's sake.

<sup>25</sup> For discussion of how responsibility for unwitting wrongdoing might derive from one's negligence, see FitzPatrick (2017).

<sup>26</sup> Alternatively, there could be instances of recklessness to trace back to, but then we can model our explanation of responsibility on recklessness.

blameworthy for causing the patient harm. How else can we explain this difference except via some tracing mechanism?

Let us grant that Joaquín is not blameworthy for harming the patient. But Joaquín also isn't negligent in treating the patient. He isn't negligent because he isn't careless; he takes appropriate precautions against causing harm. It should be plain that non-negligent actors are less blameworthy than their negligent counterparts. The only remaining question is whether explaining this difference requires an appeal to tracing. But I still don't think that it does.

Jo and Joaquín are distinguished by the fact that Jo should have known better in treating the patient, whereas it is not the case that Joaquín should have known better. But why is this true only of the former and not the latter? It might be that one's obligations are constrained by one's opportunities. Jo had an opportunity to improve her epistemic position but neglected to. Of course, so did Joaquín, at least insofar as there are steps he *could* have taken to try to find out what was in the latest issue (e.g., he could have contacted a colleague who also gets the journal). It is often the case that we could do more to try to better inform ourselves, but the mere availability of such options doesn't necessarily render our failure to take them negligent. In the specific case of physicians, the standard duty of care is not to know *every* bit of information relevant to their practice; rather, it is to know the information reasonably available to them. Indeed, the duty is to ensure, so much as is possible, that they reasonably avoid exposing their patients to avoidable risks.<sup>27</sup> Were the journal in the office (or easily readable online), and Joaquín failed to inform himself, then he might have unjustifiably run a risk of harming his patient. But he ran no such risk, since he was up-to-date, in the requisite sense, for treating the patient with due care.

The point here isn't to draw a precise boundary around negligence. I doubt one could be so drawn. Rather, it is to highlight that the question of responsibility for harming the patient just is the question of whether the physician involved was negligent, which lends credence to the idea that we can rely on whatever the standard explanation for negligence responsibility is.

## 5. Explanations and Temporal Ordering

Now, for all I've said, the tracing proponent may be unmoved. After all, what tracing proposes is that Jo is blameworthy for harming the patient only *because* she is blameworthy for neglecting to inform herself previously. This explains why she, and not Joaquín, is blameworthy for harming the patient. So while I see Jo as blameworthy because she negligently treats the patient, the tracing proponent agrees, but only because she is negligent in virtue of some *previous* episode of blameworthiness.

<sup>27</sup> Or, alternatively, that they only expose their patients to risks to which the patient consents.

The difficulty for assessing these two proposals is that verdicts of blameworthiness and negligence will naturally coincide. Return to Topsy Tina and her drunk driving, a canonical case for tracing. What makes her getting drunk wrong and her blameworthy for it? Arguably, there is nothing inherently wrong about becoming intoxicated, even to the point of being unable to control one's actions (Kantian protestations notwithstanding). Indeed, there wouldn't be anything wrong (or at least, seriously wrong) about her getting drunk if she had taken appropriate precautions. Had she taken public transportation and gone out with friends, a few of whom were entrusted as chaperones who would not be drinking, there would be little to morally criticize. Naturally, there may always be something to morally criticize about getting drunk, so long as it always impairs one's faculties, and in so doing increases the likelihood that one will act wrongly. Notice, however, that the wrong-making aspect of getting drunk here is actually inherited from *future* wrongdoing, not the other way round, as tracing proposes. It is not the case that (a) it is wrong to get drunk, (b) Tina was responsible for that and so blameworthy for getting drunk, and so (c) her blameworthiness for the accident is grounded in the prior controlled action (despite her drunkenness negating control at the time of the accident). Rather, her getting drunk increases a risk of future wrongdoing, a risk realized when she hit the car.<sup>28</sup>

The same holds true of our doctors, Rose and Jo. What is wrong about failing to keep up with medical research apart from increasing the chance that one will harm her patients? Again, it would seem, it is only because harming her patients would be wrong that each doctor's neglect to stay up-to-date is wrong. Of course, that is precisely the reason they each have an obligation to stay current in the first place. It isn't in pursuit of some noble goal of refining knowledge; it is to ensure that they're each able to provide qualified care, minimizing the risk of harm to patients.

So, there is no special reason, divorced from risking harm, for doctors to remain updated on the details of the treatments they're likely to offer. It is the future likely consequences that constrain the doctors' obligations in the cases, and which set the scope for the risk they might pose, and thus whether they are negligent or not.

Consider the case if Jo, who neglects to read the article in her office, expects to be off tomorrow. If she is called in unexpectedly, this fact seems to cut against her negligence for having failed to inform herself. She may have been planning to read the journal that afternoon. What she is obligated to do now is constrained by what is likely to be the case later. It is not negligent to forgo updating herself on the relevant research today since she is not expected to interact with patients tomorrow.

<sup>28</sup> There is, of course, an intermediary decision to drive drunk, but if she lacks control over the driving she presumably also lacks control over the decision. I take it one of the risks to take precautions against when drinking is making bad decisions, not just performing bad actions. Similarly, one might think that what is principally wrong in Tina's case is that she fails when she is sober to mitigate the risk of wrong action while drunk. But such a failure will either be reckless or negligent, and so amenable to the treatment I've given of the cases so far.

Relatedly, Rose's recklessness, when she elects not to consult the research, is dependent on the likely harms she's risking in the future. If Rose also is supposed to have the day off tomorrow, then choosing not to read today is arguably a reasonable risk to take (conditional on how reliable her schedule is, etc.). The verdicts here may be contestable, but the point is that we can modify the case so as to cast doubt on her being reckless, and we do so by adjusting the expected risk of harm.

This is partly why examples involving doctors are particularly apt. Doctors can reasonably predict fairly precisely the ways in which they will risk harm to others. But the point holds generally as well. I more or less know in my day-to-day what will be in the "scope" of my risk to others. These are just the sorts of things we talk about falling under the aegis of one's responsibilities. The proper working of my car so that it drives safely falls within that scope; bomb defusing does not. Whether I'm blameworthy for causing a harm may well then turn on whether it fell within that scope, which is just another way of saying that I'll be blameworthy when I should have known better (or, more precisely, when I ran an unjustified risk).<sup>29</sup>

Now, it's one thing to ask whether my responsibility for the accident, in a case in which I didn't take proper care of my car, depends on my failing to take proper care of the car, and another thing to say that responsibility for the accident *traces* back to it. Of course my responsibility for the accident involves my failing to take proper care of the car. I was negligent in so doing, and that fact partially explains why I was negligent in operating the vehicle, even in causing the accident. Which is just to say I ran an unreasonable risk of harm in operating that vehicle without taking proper care of it. To evaluate whether an agent has run an unreasonable risk of harm, we will of course have to look at what evidence was available to them at the time, what they knew or didn't know, what they had reason to believe, what is fairly expected of them under the circumstances, and so on. Any one of these considerations (to the extent they are separable) may be part of an explanation as to why an agent was negligent. But while these considerations may indeed be realized in the past, we should not let that temporal ordering mislead us. There is no separate explanatory step between "Jo ran an unreasonable risk of harming her patients" and "Jo is blameworthy for harming her patients" (beyond the fact that in the latter there is the actual occurrence of harm).<sup>30</sup> Of course, she is blameworthy for harming her patient because of her negligence, but this doesn't commit us to some further metaphysical relation or separate ascription of responsibility.

<sup>29</sup> The justification of risks will no doubt be partially dependent on the context of the activity. The fact that physicians' day-to-day activities routinely risk harm to patients means we'll demand more from them in terms of preventative measures than of, say, a layperson who is doing pedestrian yardwork. What is a reasonable way of conducting one's business will vary with respect to its inherent dangerousness and ubiquity. This is an element familiar in tort law, but inherited from morality.

<sup>30</sup> Importantly, Jo is negligent regardless of whether any patient in fact comes to harm, so long as she interacts with patients having not informed herself properly. But she can't be blameworthy for harming a patient if no patient was harmed.

## 6. Conclusion

Many proponents of tracing seem baffled by the prospect that it might be unnecessary. They seem unable to see “how a theory of moral responsibility could adequately handle the range of . . . cases without some sort of tracing component.”<sup>31</sup> The ambition of this chapter has been to show how we can do without tracing in addressing the epistemic dimension of responsibility. I don’t claim to have established this conclusion outright, only to have worked some way toward it. If I’m right, tracing is either unnecessary or downright problematic for explaining a range of cases central or related to culpable ignorance. We can explain these cases by extending the account I offered previously for action to target the epistemic dimensions of responsibility. Distinguishing cases based on whether the agent is reckless or negligent preserves responsibility ascriptions, but on neither score is tracing a required element of the explanation.

Proponents of tracing may insist that I’ve only redescribed what tracing is. It is important, therefore, to review what tracing explanations are supposed to be. There are cases that resist standard treatment by our theories of moral responsibility. Some agents (e.g., drunk drivers) lack the requisite control over their actions, while others (e.g., the culpably ignorant) lack the requisite knowledge of what they are doing to count as responsible on standard views. But, plausibly, drunk drivers and the culpably ignorant are still blameworthy for the harms they bring about. Enter tracing. It is meant to provide a special explanatory mechanism to capture these exceptional cases. Each kind of actor remains responsible, despite failing an ordinary condition on responsibility, because we can trace their responsibility back to a previous episode in which they did satisfy that condition (and which explains why they lack it now). Tracing is an exceptional explanation that allows us to capture aberrant cases of responsibility when our ordinary accounts get them wrong on their own.

My proposal is importantly not one that uses tracing, for the explanations I give are the very same to be given of ordinary cases. In a sense, then, what my view denies is that these are really aberrant cases. They are in a way, since they at least *seem* significantly different from, say, paradigmatic intentional action. But appearances can be deceiving. Agents can bring about outcomes in a variety of ways, even through manipulating their own mental states. These powers of agents allow them to do things in a variety of ways: on purpose while lacking important knowledge, recklessly, negligently, and besides. But these ways of doing things are already part of our ordinary thinking about moral responsibility. (Alternatively, they *should* be parts of adequate theories of responsibility.) They don’t require special explanatory mechanisms, and neither do the target cases above.

<sup>31</sup> Fischer and Tognazzini (2009: 552–3). For similar remarks, see Timpe (2011: 12) and Vargas (2005: 270).

While I have explicitly defended the view that tracing is a dispensable feature of theories of responsibility, I have not argued that it is altogether mistaken. Thus, it is open to a proponent of tracing to accept my claims here (or most of them) and insist that tracing survives. They might claim that tracing should not be characterized as an exceptional explanatory strategy, but only as the principle that responsibility at the time of action depends on past facts. Admittedly, I have no clear rebuttal to this move. But notice how strange such a position would be. For who could deny that to explain responsibility for some actions or outcomes, we will have to look at more than what is true only concurrently with the performance of that action or occurrence of that outcome. Granted, some theorists have claimed (or have been interpreted as such) that their conditions on responsibility are entirely atemporal, and so we need not look at an agent's history at all to determine their responsibility for acting.<sup>32</sup> But that is surely hyperbole, for we cannot even determine what the proper objects of responsibility should be without allowing ourselves flexibility in the size of the time-slices we examine. I cannot break the window without swinging the hammer. I cannot make a basket without moving my arms. But we (rightly) do not require that responsibility transfers from each member of that set separately to the next.<sup>33</sup>

It may be then that tracing proponents, in an effort to preserve tracing as an essential element of our theories, will wish to propose as well that tracing requires carving up actions into discrete episodes conducive to tracing's explanatory structure. If independently motivated, such a move would not be *ad hoc*, but it is a substantive commitment requiring further work than any tracing theorist has so far done.

I prefer the simpler route. Dispense with tracing altogether. All that is lost is an unnecessary special explanatory mechanism. Our theories can do without.<sup>34</sup>

## References

- Alvarez, M. and Littlejohn, C. (2017). 'When Ignorance is No Excuse.' This volume.  
 Björnsson, G. (2017). 'Explaining (Away) the Epistemic Condition on Moral Responsibility.' This volume.  
 Bratman, M. (1987). *Intention, Plans, and Practical Reason*. Harvard University Press.  
 Clarke, R. (2017). 'Ignorance, Revision, and Commonsense.' This volume.  
 Falvey, K. (2000). 'Knowledge in Intention.' *Philosophical Studies* 99: 21–44.

<sup>32</sup> Cf. Frankfurt (1987).

<sup>33</sup> To repeat a previous worry: if we did, we'd invite skeptical worries about the very possibility of responsibility; for surely we could trace back a bit further to the previous component of the action or its causal instigator, and then a bit further still. And so on, and so on. . . . Cf. Strawson (1994).

<sup>34</sup> Thanks to Neal Tognazzini, Seth Shabo, and an anonymous referee for very helpful comments on an earlier draft. For their illuminating discussion and constructive criticism I'd like to thank the workshop participants at "Responsibility: The Epistemic Dimension II" held at Vrije Universiteit Amsterdam in May 2015, as well as the organizers of that workshop and editors of this volume, Philip Robichaud and Jan Willem Wieland. Thanks as well to Josh May for his valuable comments and conversation.

- Fischer, J.M. & Ravizza, M. (1998). *Responsibility and Control: A Theory of Moral Responsibility*. Cambridge University Press.
- Fischer, J.M. & Tognazzini, N. (2009). 'The Truth About Tracing.' *Noûs* 43: 531–56.
- FitzPatrick, W. (2017). 'Unwitting Wrongdoing, Reasonable Expectations, and Blameworthiness.' This volume.
- Frankfurt, H. (1987). 'Identification and Wholeheartedness.' In F. Schoeman (ed.), *Responsibility, Character and the Emotions: New Essays in Moral Psychology*. Cambridge University Press, pp. 27–45.
- Graham, P. (2017). 'The Epistemic Condition on Moral Blameworthiness: A Theoretical Epiphenomenon.' This volume.
- Harman, E. (2017). 'When is Failure to Realize Something Exculpatory?' This volume.
- King, M. (2009). 'The Problem with Negligence.' *Social Theory and Practice* 35: 577–95.
- King, M. (2014). 'Getting Traction without Tracing: A (Partial) Solution for Control-Based Accounts of Moral Responsibility.' *European Journal of Philosophy* 22: 463–82.
- Levy, N. (2017). 'Methodological Conservatism and the Epistemic Condition.' This volume.
- McKenna, M. (2008). 'Putting the Lie on the Control Condition for Moral Responsibility.' *Philosophical Studies* 139: 29–37.
- Mele, A. (1995). *Autonomous Agents: From Self-Control to Autonomy*. Oxford University Press.
- Smith, H. (1983). 'Culpable Ignorance.' *Philosophical Review* 92: 543–71.
- Strawson, G. (1994). 'The Impossibility of Moral Responsibility.' *Philosophical Studies* 75: 5–24.
- Tenenbaum, S. (ms). 'Reconsidering Intentions.'
- Timpe, K. (2011). 'Tracing and the Epistemic Condition on Moral Responsibility.' *The Modern Schoolman* 88: 5–28.
- Vargas, M. (2005). 'The Trouble with Tracing.' *Midwest Studies in Philosophy* 29: 269–91.
- Wieland, J.W. & P. Robichaud (2017). 'Blame Transfer.' This volume.
- Yaffe, G. (2011). *Attempts: In the Philosophy of Action and the Criminal Law*. Oxford University Press.