

Traction without Tracing:

A (Partial) Solution for Control-Based Accounts of Moral Responsibility

Matt King

Forthcoming in *European Journal of Philosophy*.

This is the penultimate version. Please cite the published version, available here:
<http://onlinelibrary.wiley.com/doi/10.1111/j.1468-0378.2011.00502.x/full>

Abstract

Control-based accounts of moral responsibility face a familiar problem. There are some actions which look like obvious cases of responsibility but which appear equally obviously to lack the requisite control. Drunk-driving cases are canonical instances. The familiar solution to this problem is to appeal to *tracing*. Though the drunk driver isn't in control at the time of the crash, this is because he previously drank to excess, an action over which he did plausibly exercise the requisite control. Tracing seeks to show that an agent's responsibility for some outcome (over which he lacked control) can be traced back to a prior exercise of control which caused (in the right way) the later lack of control. These and related cases have led many theorists to treat tracing as an indispensable component of any adequate theory of responsibility. This paper argues that tracing is in fact dispensable. I offer two strategies for explaining responsibility in drunk-driving cases (and those with a similar structure): responsibility can either be exhaustively modeled on recklessness, or exhaustively modeled on negligence. Neither explanation, however, relies on tracing. If I'm right, the case for tracing is seriously weakened.

1. Tracing what

Most accounts of responsibility agree that in order to be responsible for some action or outcome, one must exercise control over that action or outcome (though they disagree about what counts as control).¹ But there are some actions which look like obvious cases of responsibility but which appear equally obviously to lack the requisite control. Drunk-driving cases are canonical instances. Suppose that Whiskey Pete is drinking in a bar and he has one (or maybe more) too many to drink. Despite being drunk, Pete gets in his car and drives toward home. In route, and due to his drunkenness, Pete hits a family sedan, seriously injuring the family inside. By hypothesis, Pete lacks the requisite control at the time of the accident. However, surely Pete is morally responsible for hitting the sedan.

This is a familiar problem for control-based accounts of moral responsibility,² but it has a familiar solution.³ To explain how Whiskey Pete could be responsible for an action despite lacking the requisite control, we can appeal

¹ Though the details can obviously differ a great deal, most theorists acknowledge that in addition to exercising a certain sort of control over their action, an agent must also meet an epistemic condition (sometimes called the knowledge condition) to be responsible for what she does. Some set it aside to concentrate on control (e.g., Fischer and Ravizza 1998), while others address it separately as raising unique puzzles (e.g., Timpe [forthcoming]; Vargas 2005), and still others locate it *within* the control condition itself (e.g., Ekstrom 2000).

² Since moral responsibility is taken to involve, principally, praise and blame, it has naturally seemed that one can only be responsible for what is under one's control. This is largely explained by the fact that something's being out of one's control appears to excuse agents from blame. Thus, what amounts to the control required to be morally responsible has been the primary inquiry for theories of responsibility. An extremely popular account holds that what matters is 'responsiveness to reasons' (e.g., Fischer & Ravizza 1998; Wallace 1994), while others claim it is indeterministically caused decisions or choices (e.g., Kane 1996),

to *tracing*. Tracing is a special explanatory tool that denotes a backwards-looking strategy for explaining ascriptions of responsibility precisely in cases which lack the requisite control. We can still vindicate ascriptions of responsibility without control if we can ‘trace back’ to some prior exercise of control suitably connected to the current conduct. If we look at Pete’s past, we notice that he drank to excess. And this was an action over which he did plausibly exercise the relevant control. If Pete is responsible for getting drunk, and his drunkenness caused the crash, then these facts could explain why Pete is responsible for hitting the sedan.⁴ Tracing seeks to show that an agent’s responsibility for some outcome (over which he lacked control) can be traced back to a prior exercise of control which caused (in the right way) the later lack of control.

It is important to emphasize that explanation by tracing is an *exceptional* explanation, meant to handle exceptional cases. While we can normally explain responsibility by appeal to exercises of control at the time of action, these cases resist such treatment because the usual control is absent. So we’re forced to look to the agent’s past. If Paul coolly decides to run into his ex-wife’s sedan in retaliation for some perceived wrong, he plausibly exercises control over the entire action. We needn’t look to the past at all to explain his responsibility for the crash; he steers the car intentionally, with full awareness of his actions, etc. Not so with Pete. He is so drunk as to not be in full control of his current actions, but the reason for his reduced capacity lies in his past (controlled) actions. In this way, Pete’s case reflects the canonical schema for tracing:

- (1) The agent A did something x without control at t_1 ;
- (2) A performed action y with control at t_0 ;
- (3) A’s action y at t_0 made A have property P; and,
- (4) P explains why A lacked control at t_1 and why A did x.

Pete hit the sedan at t_1 , where he failed to control his action. But Pete drank to excess at t_0 and did so while satisfying the control condition (stipulated). Moreover, his drinking to excess made him drunk, and his drunkenness explains why he lacked the requisite control at the time of the crash (and, indeed, why he crashed). Tracing is most readily adopted in cases like Pete’s, where an agent has done something in the past to change his or her condition such that in the present the agent does not meet the control condition on responsibility.

One of the most striking aspects of tracing is its popularity. The vast majority of views seem to hold that tracing is necessary to explain at least some important cases of responsibility. Whatever other differences between theorists, there seems widespread consensus that tracing is a necessary explanatory tool.⁵ Tracing’s ubiquity has even led some to claim that tracing is ‘theoretically indispensable’ (Fischer & Tognazzini 2009: 40); a theory of

while still others find the required control in autonomous action (e.g., Mele 1995, 2006), and yet others require agent-causation (e.g., O’Connor 2000). Some have even assumed control necessary as part of skeptical arguments about moral responsibility (Strawson 1994). Despite the differences among these views, they all agree that uncontrolled action is not the proper target for our moral evaluations, and so agree that control is required for responsibility. (Of course, not every theorist agrees that control is required; see n.24 and n.28).

³ Both Fischer and Ravizza 1998 and Mele 1995 contain considerable defenses of a solution, though their targets are more diverse than just drunk driving cases. McKenna 2008 and Vargas 2005 provide good overviews of the literature, in its varied forms, along with discussion of some complications.

⁴ We may also have to include a condition that Pete knows that drinking can impair one’s driving and increase the risk of harm to others, but this should be a relatively uncontroversial assumption. In general, the epistemic condition requires that one bear the right epistemic relation to that over which one exercises control. In the case of tracing, it looks like one may need to bear that relation over the action for which one is supposedly responsible. For more discussion of the epistemic criteria on responsibility and tracing, see Fischer & Tognazzini 2009, Timpe (forthcoming), and Vargas 2005. I’ll say a bit more about foreknowledge as well in what follows.

⁵ As Manuel Vargas has noted, ‘[tracing] is one of a few things to which nearly all parties in the debate about free will appeal to with equal enthusiasm’ (2005: 270).

responsibility could not adequately explain an important range of cases without reliance on tracing. How else can we capture Pete's responsibility for the crash?

If tracing is indeed indispensable, then we should want it to be free of problems of its own. But recently tracing has been the subject of a number of critical discussions (McKenna 2008; Sher 2009; Smith 2008; Vargas 2005). In different ways, theorists have suggested that tracing is not unproblematic after all. If tracing were obviously necessary, one might view these critiques as merely suggesting puzzles for future research, rather than presenting significant obstacles for theories of responsibility. After all, if everyone must appeal to tracing, then whatever problems tracing presents are everybody's problems.

This paper argues that tracing is, in fact, theoretically *dispensable*. I sketch a strategy for handling drunk-driving cases without resort to any significant notion of tracing. My goal will be to illustrate that there is a defensible view of drunk-driving cases (and any case that shares the structure of such examples) that does not require appeal to tracing. Since these cases are considered canonical examples requiring tracing, if my argument is successful, the case for the necessity of tracing is weakened.

Moreover, if I'm right and tracing is not in fact necessary, the recent critiques of tracing become problems only for views that appeal to tracing. They must be addressed, not as puzzles for all views to work out, but as a significant obstacle to only certain theories of moral responsibility. (Additionally, those not already wedded to the notion of tracing need not worry about whatever difficulties tracing may invite, for there is a plausible alternative account.)

My argument for the claim that tracing is not in fact necessary to explain responsibility in drunk-driving cases centers on denying that we need a special explanatory tool for such cases. I argue for an alternative explanation, modeled on perfectly ordinary cases of action. The upshot is that if we can jettison tracing from our accounts, we can avoid the special difficulties tracing may invite, potential problems facing virtually every orthodox view.

Before I turn to discussing drunk-driving cases, a brief note is in order. As I've said, tracing is canonically associated with drunk-driving cases. But it is the focus as well in so-called manipulation cases.⁶ In such cases, the agent seemingly exercises control over his action, but appears *not* to be responsible in virtue of having been psychologically manipulated so as to perform the action he does. Here, many theorists argue that what explains the agent's non-responsibility is that we can 'trace' backward to the manipulation. If tracing is required to avoid getting the wrong result in manipulation cases, then, even if I'm right about drunk-driving cases, tracing might still be theoretically necessary. I think, however, that tracing is being put to quite different purposes in such cases, though I do not have the space to consider the point here. The argument that manipulation cases also do not require an appeal to tracing must be considered elsewhere.⁷

2. Drunk-driving cases and the planfulness of agents

Drunk-driving cases are plausibly paradigm instances requiring tracing. How else can we explain responsibility in Whiskey Pete's case, other than by appeal to tracing? Indeed, for some theorists, cases like Pete's may be the only

⁶ See, for example, Fischer, et al 2007: 93-98,211-215; Frankfurt 2002; Haji 2008: Ch 6; McKenna 2008; Mele 1995, 2006; and Pereboom 2001.

⁷ Tracing has also been used as part of an account of responsibility for one's emotions or character traits. See, for instance, Fischer & Ravizza 1998: 255. My aim in this paper is responsibility for actions or outcomes and I mean to take no stand on how to account for responsibility for emotions or attitudes. It is perhaps worth noting, however, that whereas control is a standard requirement for responsibility for actions, it is far more contentious as a condition on responsibility for attitudes (see Sher 2006; Smith 2008).

cases in which an agent's history matters. But I think that explaining drunk-driving cases does not require tracing at all. My argument offers two viable strategies, assuming for its purposes that Pete is indeed responsible for the crash. I'll argue that Pete's responsibility for the crash can be either (i) exhaustively modeled on cases of recklessness; or else (ii) exhaustively modeled on cases of negligence. Which of these strategies one takes will depend upon what one thinks about Pete's psychological states leading up to his leaving the bar. But I will argue that one or the other of the above explanations will suffice for explaining Pete's responsibility, but neither of them requires tracing. So Pete's responsibility for the crash alone does not push us toward accepting tracing as an explanation for why he is responsible. At the very least, if my argument succeeds, we will have a plausible alternative to using tracing.

Suppose a soldier wants to incapacitate a group of enemies hunkered down in a pillbox. So, the soldier lobs a grenade into the pillbox, incapacitating the enemies. We can imagine that for several seconds after he has thrown the grenade, indeed, for a few moments after the grenade has landed in the pillbox, he still hasn't achieved his goal. Only when the grenade finally explodes does he incapacitate them.⁸ Similarly, we can imagine a different soldier with the same goal, incapacitating a group of enemies, but this time he is being pursued by them. So, the soldier sets a mine along his path, intending for his pursuers to trip it, and thus be incapacitated. Sure enough, his ploy works. However, the mine isn't triggered for several minutes after he sets it. He is a good half-mile away when it finally goes off and the enemies are incapacitated. A final example: suppose a yet different soldier is planning to rescue some captured comrades. His plan involves incapacitating several groups of soldiers who are positioned among the guardhouses and towers. So, he goes about setting explosives at each structure, timed such that they go off all at once. He is a safe distance away when they do explode, and he successfully incapacitates all the enemies. But it has been nearly an hour since he set the first charge (caution required he take his time).

These three examples are of interest here for three features they each possess. First, I simply assume it uncontroversial that all three soldiers are each responsible for incapacitating enemies. So, the first feature is that the agent involved is responsible. Second, I take it as obvious that tracing isn't needed here. However we construe the requisite control condition, the soldiers must satisfy it. After all, these are paradigmatic enough instances of intentional action. Barring skeptical views about responsibility in general, there's little reason to deny responsibility, and therefore to deny control, to their actions. But the final feature of interest here concerns how each soldier exercises control. All three cases illustrate how agents can control outcomes despite being removed, physically and temporally, from the event's completion. The grenade example shows how an agent's contribution is typically exhausted a bit prior to an event's completion, even for actions that are accomplished quite quickly. But as the cases progress, what the soldier does in incapacitating the enemies is further and further removed, temporally, from the event of their being incapacitated. Nevertheless, this time lapse does nothing to modify the account we need to give in order to explain the soldiers' responsibility, nor does the fact that each soldier fails to exert any direct physical control over the outcome at the moment it occurs. Each has a goal (incapacitating the enemies), he constructs a plan to achieve the goal, and he executes the plan. In the latter cases, this means setting things up such that later events occur (e.g., the enemies' tripping of the mine; the explosion of the first charge).⁹ We don't need to appeal to a special tool, like tracing, to explain responsibility in such cases. The soldiers use the ordinary planning capacity of agents to perform regular actions. These are typical cases, in structure if not in content, not marginal cases requiring special explanation.¹⁰ In all of the above cases the soldier is clearly responsible for incapacitating the enemies, an action over

⁸ And even then there is a small time lag between explosion and incapacitation.

⁹ Even in the first case (with the grenade) this is true. He presumably pulls the pin to arm the grenade, moves his arm to get the grenade in the pillbox, and the grenade explodes. Each of these events is important to realizing the goal of incapacitating the enemies.

¹⁰ For related cases, see Sergio Tenenbaum's (forthcoming) discussion of 'gappy' actions. These actions persist across time (like taking a walk) though they may be interrupted by other actions (admiring flowers) without one's ceasing the original action. Such individuals plausibly control their taking a walk even while stopping to smell the roses.

which they seemingly exercise control. Yet they all are (to various degrees) removed from where the action is when the outcome actually obtains.

To emphasize: it isn't that we must trace responsibility for incapacitating the enemies to some previous act the soldier is responsible for (e.g., pulling the pin to lobbing the grenade; setting charges 1-4). It's rather that his contribution to bringing about the terminal event is completed before (sometimes long before) that event occurs. Nevertheless, the soldier's intention or plan is enough to explain responsibility for the subsequent terminating event. On the assumption that control is required for responsibility,¹¹ a natural inference to draw is that the soldiers must therefore exercise control over the terminating event.

As Bratman (1987) has noted, agents have directed attitudes which organize their conduct over time.¹² Whatever particular account we give of this capacity, all that is needed here is the basic observation. Intentional action with full awareness of the likely (and intended) consequences is a paradigm instance of responsibility for those consequences. And as a paradigm, explaining responsibility in those circumstances requires no special explanatory tool, be it tracing or something else. Indeed, it seems sufficient of the soldiers' cases to say that each intentionally set a plan into motion that reached fruition with the incapacitation of the enemies. Each satisfied the control condition in constructing and executing the plan, even though they did not exercise direct (physical) control when the plan succeeded. (Note that none could have prevented the incapacitation of the enemies moments before it occurred.) These must be standard cases of action, thus they seemingly call for a standard explanation of responsibility. Tracing is unnecessary here.

Now, all the soldier examples involve intentional action. The soldier is trying to realize a particular end. But Whiskey Pete's case isn't like this. Pete doesn't intend to hurt the passengers in the sedan; he doesn't have a plan. Instead, Pete just knew or should have foreseen that his getting drunk and then driving would put others in serious risk of harm, that he would be much more likely to injure people.

So Pete's case is somewhat different from the soldiers. Nevertheless, we can draw the relevant parallel here by considering a soldier who lays a mine to incapacitate his pursuers, only to have it detonated by a child from a nearby village who was wandering through the forest. As in Pete's case, the soldier here doesn't intend to do any harm to the child; it isn't a component of his plan. In both these sorts of cases, then, the first step involves distinguishing between two possible interpretations of Pete's case. What matters here is whether Pete actually foresaw that getting drunk might lead to him injuring others on the roads or, instead, whether he *should* have foreseen that this was the case (but didn't). If he actually considers the possibility of future harm at the time of his continued drinking, then his case can be usefully modeled on recklessness, which I shall argue involves exercising control in much the same way as the soldiers' cases above (involving intentional action).¹³ Call this the reckless interpretation.

If Pete doesn't consider the possibility of harm, but should have, then we can model his case on cases of negligence. Call this the negligence interpretation. I will argue that tracing is unnecessary here, too, for tracing cannot explain responsibility in negligence cases. The next two sections consider the recklessness interpretation and

¹¹ If control is not necessary for responsibility, then tracing isn't a viable mechanism for explaining Pete's case. But, then again, if control isn't necessary, drunk-driving cases don't present a puzzle, since Pete's lack of control at the time of the accident won't thereby be an obstacle to an ascription of responsibility. Thus, the assumption of a control condition is benign with respect to debate about tracing.

¹² Here I am relying only on Bratman's general observations about the structuring nature of agency, not his account in all its details.

¹³ Steinbock 1985 argues that drunk driving often involves recklessness, but focuses on legal concerns. Her argument seeks to show that drunk drivers exhibit the requisite malice to justify charging them with murder. She is explicitly interested in neither control nor its relevance to tracing.

the negligence interpretation, respectively. On neither interpretation, however, is tracing needed to explain Pete's responsibility for the crash.

3. Tracing and recklessness

If Pete actually thinks about the possibility of his harming someone due to his getting drunk, then the mechanism for explaining his responsibility for the crash can mirror the explanation of the soldier's responsibility. Though he did not intend or plan on harming others, I claim he satisfied the basic control needed for responsibility. By considering the possibility of harm and ignoring it, Pete acts in a reckless fashion. He is aware that continued drinking will leave him impaired, and that he might injure someone as a result. But should he injure someone, (as he in fact does), we need not 'trace back' to his responsibility for continuing to drink. If the soldier controls incapacitating the enemies when the charges he set go off, it is plausible to suppose that Pete exercises the requisite control over the possible, considered harms. There is no need to trace backward anymore than in ordinary action; the structure required extends control forward even when Pete is no longer able to directly affect the course of the car.

As an initial defense of this claim, that Pete is responsible in virtue of exercising control, consider that reckless results (especially harms), while not intentional, still count as having been done knowingly. Pete can't excuse himself by claiming that he didn't mean to harm the sedan passengers. Of course, this is true on a strict reading of intentionality. But what is truer still is that Pete didn't mean to not harm anyone, for if he had, then consideration of the harm he was risking by continuing to drink (when he would likely drive) would have led him to stop. But reckless actions aren't actions for which the requisite control is absent. I take it that reckless harm is an ordinary enough instance of action. Rick thinks to himself that if he exceeds the speed limit on slick roads, swiftly switching lanes, he might well cause an accident. But, so be it, he does so anyway. Sure enough, he causes an accident, hitting a sedan and injuring those inside. While not a paradigmatic case of *intentional* action, Rick's case isn't far from central cases of ordinary action. Again, on the assumption that responsibility requires control, and since Rick is surely responsible, he plausibly satisfies the control condition over the accident. If Rick is responsible in virtue of exercising control, and if Pete's case mirrors Rick's, then I conclude that Pete, too, is responsible in virtue of exercising the same control.

One might seek to resist concluding that Pete is responsible in virtue of exercising control. One way to do this would be to insist that while the cases are parallel, they both actually involve tracing. The claim would have to be that Rick's responsibility doesn't hinge on exercising control over the accident, but rather that we should trace back to his exercise of control over the decision to drive in a dangerous fashion. This would preserve the similarity to Pete's case, where the tracing proponent advises tracing back to Pete's control over continuing to drink.

But this is an unpromising strategy for two reasons. First, it is just too implausible to suppose that tracing is called for in Rick's case. We have little independent reason to suppose that he doesn't exercise sufficient control over his driving even at the time of the accident. While in Pete's case his degree of drunkenness plausibly incapacitates his control, Rick just drives a bit dangerously. Since we have no reason to suppose his control impaired, and since he is uncontroversially responsible, tracing looks superfluous here.

Second, and more importantly, to extend tracing to cover cases of reckless action would be to significantly revise how tracing is commonly understood. Recall that I've characterized tracing as an exceptional explanation to cover exceptional cases. Whatever else one might say about reckless action, it isn't particularly exceptional. We knowingly adopt risks when we decide to drive a bit dangerously, but also in pursuing hobbies (like skydiving) or investing money or opting not to repair a small roof leak. If tracing becomes the mechanism for explaining even reckless action, then it is no longer a special explanatory tool, but rather a regular part of our ordinary account of moral responsibility. As such, claims that tracing is 'theoretically indispensable' become somewhat trivial.

For all I've said, then, one potential response to my argument here is to conclude that tracing is ubiquitous. Every action, from the soldiers' fully intentional plans to semi-conscious drunk-driving will necessarily involve tracing. (This follows from the structure of my argument, which is to draw a parallel between canonical tracing cases and ordinary action.) This would still be a remarkable result, however, since the prevailing view in the literature is that tracing is needed to explain a certain *range* of cases, not *all* cases. Moreover, this result would seemingly undermine the primary *motivation* for an appeal to tracing. To get the right result in cases like drunk-driving, so the thought goes, we must employ tracing. If tracing is to be used everywhere, this seems to weaken its initial motivation.

I take it, therefore, that tracing proponents aren't likely to adopt this first strategy. The other way to resist Pete's responsibility by virtue of control is to claim that his case isn't parallel to Rick's. While Rick is under control throughout his reckless driving, Pete not only ceases to control his driving once drunk, he ceases to exercise control over any of his actions. Thus, the tracing proponent contends, Pete's drunkenness impairs his capacities of control in a way that Rick's capacities are not similarly blunted. So while Rick's assumption of the risk involved may be sufficient to extend his control over his subsequent driving, given his unimpaired later mental states, Pete's assumption of the risk does not extend control to his later driving. Tracing is thus required for Pete after all.

The promise of this strategy rests on undermining Pete's case as a case of recklessness. We must conclude that while Rick's knowingly hitting the sedan suffices to imply control, Pete must not knowingly hit the sedan in virtue of his impaired condition. I cannot convincingly show here that this could not be a viable strategy. But I do think we have good reason to suppose that Pete still knowingly injures the sedan passengers, despite his drunkenness. After all, he can't very well claim that he didn't think anyone could be injured or that he was doing anything dangerous, even though, at the time of his driving, he was perhaps numbed (or even oblivious) to the risk he was posing.

Consider the soldiers again. Even though the soldier may regret laying the mine, or think that he should have taken a different course of action, the plan has been set in motion, and there is nothing he can now do to alter its course.¹⁴ It would strain usage, I think, for the regretful soldier to say he didn't mean to incapacitate the enemy, even excepting the strict intentional reading. Relatedly, while he can say he didn't mean to incapacitate the child, this only works as a denial of the result being intended. He can't equally lay claim to having meant to *not* hurt the child, since he took no precautions to avoid that harm. This is partly what makes him reckless. If this is the right thing to say about the soldier's case (and I think it is), then recklessness may not require maintaining the awareness of risk at all. Indeed, should Rick become distracted by the radio, and forget about the risk he's running (or the soldier forget laying the mine), it still seems as though he knowingly drives dangerously, though he is now unaware of the risk.¹⁵

To illustrate the point more vividly, suppose Jill pops a pill she knows makes her likely to become an 'unconscious zombie' who suffers random and serious muscle spasms 30 seconds after taking it. She then walks into a china shop. While she certainly lacks the requisite control for the duration of the pill's effect, she is also certainly responsible for all the china she breaks due to her spasms. But this needn't be because she is responsible for some prior action, the taking of the pill, which is causally related to the absence of certain mental states which led to the

¹⁴ It is interesting to note that when agents change their plans this could have a bearing on their blameworthiness or praiseworthiness even if it didn't affect the action's intentionality. Regretting a decision or changing one's stance toward a goal, even when one can no longer avert that outcome, can still have a bearing on one's moral evaluation.

¹⁵ Perhaps a related objection would insist that Rick's recklessness lies in his driving decisions, which themselves cause the crash, whereas Pete's recklessness lies in his drinking decisions, which themselves do not cause his crash. Thus, Pete's case, even if understood as a case of recklessness, isn't strictly analogous to Rick's. I don't think this objection succeeds. Rick's recklessness may concern his driving exceptionally quickly in especially wet road conditions. This recklessness sets background conditions for the crash, but the crash is caused by, say, his failure to be able to stop in time or adjust his course. In the same way, Pete's drinking sets the conditions that make the crash possible, such that he doesn't have time to stop or alter his course. The parallel between the two cases remains sufficiently intact.

later harm. It is rather that she is already related to the later harm at the time she takes the pill. Her foreknowledge of its effects coupled with her beliefs about the shop and the likely harm she would cause in it are sufficient for explaining her responsibility. She acts recklessly by entering the shop after having taken the pill, knowing as she does the risk of harm she'll be running. She clearly breaks the china knowingly, even though she doesn't know (at the time of its breaking) that she is breaking it. No tracing mechanism is needed.¹⁶ So if Jill is responsible, so too, is Pete.

To emphasize again: Jill realizes that in taking the pill and walking into the shop she substantially increases the risk that she will break some china. In carrying through on that action, should she break china as a result, she breaks it knowingly, though she is no longer in direct control of her actions at the time. To explain her responsibility it is sufficient to point out that she breaks it knowingly. And if this is sufficient for explaining her responsibility, which requires control, then she seemingly exercises sufficient control over the breaking of the china. This explanation is, of course, compatible with there being ways she could have exercised increased control over the breaking of the china. Should she have not taken the pill, and instead entered the shop with a baseball bat and, with clinical precision, obliterated setting after setting of fine china, she would also be responsible, and would have exercised increased control over her actions. After all, by monitoring her conduct, she retains the ability to fine-tune her movements, the specific sub-actions and manner of performing them which constitute the entirety of 'breaking the china'. But there can be little doubt that there are degrees of control which outstrip the threshold required for responsibility. I need not be an expert marksman to sufficiently control the shooting of my landlord, and thus be responsible for it.

One might object that Jill's exercising control over the breaking of the china is incompatible with the fact that she is, simply put, out of control at the time the china breaks. The pill renders her a non-agent, so we must trace backward to find the necessary control. But this doesn't seem quite right. If Jill were to intentionally take the pill with the intention of breaking the china, the fact that she was oblivious (at the time) to the risk of harm she posed would be of little relevance to her intentionally breaking the china. There are many ways that agents can do things. If Jill wants the china broken, she could hire someone to do it for her, rig a complicated machine and set it loose in the shop, arrange for a traffic accident that sent a truck careening through the plate window, etc.¹⁷ Or, she could take a pill that makes her likely to flail wildly, and then enter the shop. The fact that her agency is compromised, that she loses the ability to directly control her actions, after the pill kicks in, is really an orthogonal feature of the case. She, quite craftily, uses *herself* as an indirect instrument of destruction. Nevertheless, she breaks the china knowingly. And, if the soldiers, Rick, and all manner of temporally-extended action-performing agents exercise sufficient control to be responsible, then Jill does as well. And if Jill is responsible in virtue of exercising control, so is Pete.¹⁸

Moreover, Jill's case isn't unique in structure. We can draw up additional cases wherein agents plan to achieve or foresee achieving certain ends precisely by numbing themselves to the current circumstances (e.g., a soldier in combat, a heroic rescue worker), or by otherwise removing vital mental elements. In doing so, they don't present a special obstacle to ascriptions of responsibility when they act, one requiring an exceptional explanatory

¹⁶ It should be noted that this explanation is perfectly compatible with the possibility that Jill doesn't break any china in the shop. Perhaps she takes the pill, it works, and yet, in a strange turn of events, her flailings never break any china. The fact that in being aware of the risk of harm she is thereby responsible for the harm is not undermined in cases where that harm isn't actually realized. In such cases she is not responsible for breaking the china for the simple reason that no china is broken.

¹⁷ Whether or not these all would count as instances of Jill *intentionally* breaking the china doesn't matter here. All that needs to be shown is that agents can do things for which they are responsible (and without appeal to special explanatory tools like tracing) in a wide variety of ways.

¹⁸ It may be that, as one referee suggested, we react differently to cases like Jill's wherein the agent lacks direct control over her bodily movements, because we ordinarily have direct control over those movements. Since it is those movements which cause the harm in question (unlike the soldiers' cases), this makes at least an intuitive difference. I'm inclined to think this is right, though I take no stand on where the intuitions originate. But the existence of an 'odd' way of controlling one's bodily movements isn't enough to suggest that it requires an exceptional explanation. Part of the creativity of agents is their ability to do ordinary things in extraordinary ways.

tool. They merely use their ordinary capacities as agents. Indeed, this general explanation is common enough for responsibility in many cases of everyday action. I place weedkiller on my lawn knowing that should it rain (which is it likely to do) it will wash into my neighbor's prized petunia patch, killing her beloved flowers. I need no special tracing mechanism to explain my responsibility for killing her flowers when it does in fact rain. I've acted recklessly so long as the chance of rain was sufficiently high. Or take any case of non-intended consequences brought about through some intentional action. It simply seems superfluous to rely upon tracing for such cases.

It's worth reminding ourselves that while it is open to the tracing proponent to try to capture all these cases of knowingly doing things by appeal to tracing, this will seriously revise the notion of tracing. They can claim that Jill breaks the china knowingly only because there were actions in the past, her taking of the pill and entering the shop, over which she had control. But as I argued above, to say that all actions require tracing simply looks implausible for ordinary and paradigmatic actions, and seems to undermine the principal motivation for an appeal to tracing in the first place.

Alternatively, one might think that the difference between intending and foreseeing an outcome is an important difference between the cases. While there is obviously a difference between intentional outcomes and foreseen outcomes, the distinction and its relevance to moral responsibility are orthogonal to the discussion of tracing. For my argument merely concerns the claim that tracing is a special explanatory tool reserved for cases wherein agents seem to fail to satisfy the paradigmatic conditions on responsibility. My claim has been that canonical cases of tracing can be explained in exactly the same terms as more 'ordinary' cases of action.

Indeed, the distinction between intending and foreseeing helps support this point. If one thinks that agents are less responsible for merely foreseen harms than for wholly intentional ones, then one ought to conclude that Pete and Jill and Rick are less responsible than the soldier. If one thinks there is no difference between foreseen and intentional results, then one should conclude they are all equally responsible. But no matter which conclusion one reaches, one will reach it in virtue of one's stance on the question of the intentionality of an action on an agent's responsibility, not whether tracing is an exceptional tool. Consider an agent who merely foresees that his pushing a button will incapacitate his neighbors. He may well be less responsible than the soldier, but he isn't obviously less (nor more) responsible than Jill or Pete. So while the distinction between foreseen and intentional outcomes is no doubt relevant to discussions about moral responsibility, it doesn't seem to bear on the significance of tracing itself.

In all of the above cases, an agent brings about some end intentionally or knowingly even when they no longer directly control their current actions or the state of the intended result. And in every case we can claim that they exercised sufficient control over the resulting event, even when that event occurs much later. If Pete thinks about the potential consequences of his getting drunk when he continues to drink, I conclude that this fact, too, is capable of helping to explain his responsibility for the later harm, for it shows that harm to be done knowingly and for Pete to have acted recklessly. And if drunk-driving cases can be explained via a mechanism modeled on recklessness, then tracing is not necessary for explaining responsibility in such cases. The remaining option for the tracing proponent is to significantly revise the notion of tracing, from an exceptional explanation of problematic cases to a standard (and perfectly ordinary) mechanism of agency. Neither result looks promising for the proponents of tracing.¹⁹

¹⁹ It should be noted that no proponent of tracing takes it be part of standard explanations of responsibility. It is always meant to be a special tool used for difficult or exceptional cases. And should a tracing proponent wish instead to argue that tracing is ubiquitous, it seems she will need novel motivation for the position, for she will not be able to work from 'exceptional' cases like drunk-driving.

4. Tracing and negligence

But suppose Whiskey Pete just doesn't think about the consequences. He's having too good a time at the bar, and while he might agree if asked that of course alcohol impairs driving ability and increases the risk of harm, this thought never actually occurs to him. We might then say that he *should* have realized that by getting drunk he was increasing the likelihood that he would injure someone by driving home. This is something he knows, something any competent adult knows, and it should occur to any competent adult when one decides to continue drinking.

One who has these thoughts is thinking that Pete is negligent. He fails to consider the increased risk of harm he will pose to others if he gets drunk and then drives. He fails to take due care to limit the risk his conduct poses. Similarly, if Norma is distracted while backing up out of her driveway and fails to check her mirrors, and as a result she hits a child, she is responsible for hitting the child because she was negligent. Norma fails to take due care to limit the risk of harm her backing out poses to those around her by failing to look behind her. She is careless, as Pete is careless in not thinking to take certain precautions to avoid posing an unreasonable risk of harm to those around him (either by refraining from additional drinking or some other preventative measure, like giving up his keys).

What makes Pete and Norma's cases similar is that neither one needs to have thought about the risk of harm. Both of them are competent adults, and so 'know' in some dispositional sense the dangers of driving carelessly. But negligence doesn't require that they actually consider this risk. Pete can be wholly unconcerned with what might happen later that night as he orders another round; Norma can be absorbed in a song on the radio or thoughts about her difficult meetings later that day. It will suffice for negligence that each should have done more to limit the risk of harm, and yet they failed to do so. Failing to take the reasonable measures to curtail the risk of harm one's conduct poses is sufficient for one's being negligent.

Intuitively, I think, we regard both Whiskey Pete and Negligent Norma as being responsible. If Norma is responsible, then it is because she is negligent. And if I've given an adequate characterization of negligence, the same explanation can be given of Pete's case. Pete displays a similar carelessness to Norma's. So if Norma is responsible for hitting the child, then negligence can explain Pete's responsibility for the car crash in the same way.

Just *how* do we explain responsibility in Norma's case? Fortunately, for my purposes here, the details of the explanation don't matter. I will remain neutral on just how we explain why negligent agents are responsible for the products of their negligence. Instead, I will show that, whatever the explanation is, it will not need to involve tracing; indeed, I don't think explanations of negligence cases *can* rely on tracing.²⁰ Thus, if tracing cannot explain responsibility in cases of negligence, then on the negligence interpretation, tracing cannot explain Pete's case either, and so not only is tracing not necessary for explaining drunk-driving cases, it isn't even sufficient.

First, it's important to realize how tracing is supposed to work in drunk-driving cases like Pete's. We're supposed to trace back from Pete's hitting the sedan to his decision to continue drinking. Over that decision he supposedly did have the requisite control, he is responsible for it, *and* it put him in the condition of impairment that led to the crash. But such conditions don't apply in Norma's case. She, we may suppose, didn't do anything for which she's responsible that put her in the condition of distraction. She was just distracted by the radio (which was left on by the previous driver), or by her upcoming meetings (which simply 'popped' into her head). So if we are going to explain Pete's responsibility by modeling our account on Norma's case, tracing won't need to figure into it.

One might object that the above simply makes the argument for treating Pete's case differently than Norma's. But this would simply take us back to giving Pete's case the foreknowledge interpretation discussed in the previous section. Perhaps that's the better (or more felicitous) interpretation of these cases, but I take no stand on that question here. Instead, I simply note that if Pete's case is to be treated like standard cases of negligence (e.g.,

²⁰ This isn't to say tracing hasn't been employed in attempts to explain negligence responsibility. R. Jay Wallace argues that we must trace back to earlier episodes of choice to blame the negligent (1994: 139), whereas Michael Zimmerman argues we must trace back to an actual instance of considering the risk posed by the conduct (1986). The argument here is meant to illustrate why such attempts fail. I defend a similar line of reasoning (in greater detail) in King 2009.

Norma's), then tracing can't do the work. He will simply be responsible for the accident for the same reasons that Norma is responsible, because both should have taken better precautions to limit the risk they posed to others.

Proponents of tracing might opt for a different tack. They might claim that we ought to explain Norma's responsibility (and Pete's) by tracing back in time to some opportunity each had to exercise a capacity to be non-negligent but omitted exercising that capacity. But omissions can be understood in two ways. First, one can choose (understood loosely) to omit to do something. In the relevant cases, this would mean acting without doing something one knows one ought to do. Suppose Pete chooses to continue drinking in the awareness that without considering the consequences his action is likely to lead to bad results. This is reckless behavior, and can be explained via the mechanism discussed in the last section concerning doing things knowingly.

Second, omissions can be simple failures to do something. In the relevant cases, this would mean acting without doing something one ought to do, without realizing that one ought to do it. On this view, however, a regress threatens. Such past omissions would constitute *further* instances of negligence, and therefore cannot make headway at explaining negligence itself. Tracing in this instance would merely push our explanation backwards to some previous instance of negligence, responsibility for which would itself still require explaining.

So far, I've argued that should Pete not think at all about the potential consequences of his continued drinking, then he acts negligently, and his responsibility for the crash can be explained by whatever mechanism explains responsibility for negligent behavior generally. I take no stand on what particular mechanism that is; however, I've argued that whatever the explanation, tracing cannot figure into it. Either there is nothing to 'trace' back to in cases of negligence, or else the relevant omission would itself be an instance of negligence, and so would be of no explanatory help.

One might here object that I've mischaracterized negligence.²¹ Norma's negligence isn't in her hitting the child, it is characterized by her failing to take sufficient care in backing up (by, say, checking her mirrors). Explaining her responsibility for hitting the child is then traced back to something she is responsible for, namely, her failure to check her mirrors. We are not, therefore, explaining negligence via negligence, we are merely explaining her responsibility for the injury by reference to her negligence, which she is responsible for. But while I have remained neutral with respect to how negligence explains responsibility, the objection cannot maintain the same neutrality. This is because in order for tracing to succeed, the account of responsibility for negligence must give an explanation of the negligence that doesn't also apply to the resultant outcome. For if the explanation can be directly applied to the hitting of the child, tracing again is unnecessary, since we can give a direct account of Norma's responsibility for the injury itself.

Unfortunately, attempts of this sort to account for Norma's negligence will likely eliminate the regress only by eliminating the need for tracing. For on a plausible general account of what counts as negligence, that account will be reproducible with respect to Norma's hitting of the child. Consider a reasonable candidate which holds that Norma is responsible for failing to check her mirrors because she had it in her powers to check them but failed to do so as a result of a normally-operating mechanism. This condition also seemingly applies to Norma's failure to avoid hitting the child. It, too, was in her power to avoid hitting the child and she failed to as a result of a normally-operating mechanism (at least on a plausible understanding of normal operation). And if the explanation used to account for Norma's failure adequately explains her responsibility for that failure, and if it can be applied to her hitting of the child, then there's again no need to trace back to her failure in order to explain her responsibility for hitting the child. If the explanation serves for the prior failure, it should serve for the later injury. So even on the objection's view, tracing is unnecessary. For we'll be able to directly account for the injury by some negligence explanation. And this result isn't unique to the particular accounting of negligence considered here. However the failure to take reasonable precautions to avoid causing harm is explained it will be extendable to the failure to avoid

²¹ My thanks to Neal Tognazzini for suggesting this line of thought.

actually causing that harm, at least where general negligence is concerned, precisely because negligence is characterized by the risking of harm without taking due care to avoid it. And this characterization is applicable both to the causings of harm with which we're concerned as well as prior failures to take reasonable precautions. Both kinds of failures unreasonably increase the risk of harm.

Granted, her failing to avoid hitting the child is explained by her failing to check her mirrors. But the objection cannot press this point too strongly, because this failure, too, may be explained by previous failures (e.g., to have her morning coffee, or set her alarm early enough). Thus, the regress threatens. So, either the regress is eliminated and the need for tracing along with it, or else one may attempt to salvage the significance of tracing only by inviting a problematic regress, wherein we have no grounding for the subsequent responsibility to trace back to.

A final objection might seek to drive a wedge between Norma's case and Pete's. Perhaps we cannot adequately isolate the negligence mechanism with respect to Norma's failure to check her mirrors., and so such an explanation will apply to her hitting of the child as well. But this is because she had it in her power avoid hitting the child, since she had it in her power to check her mirrors. This is not equally true of Pete, who, due to his drunkenness, did not have it in his power to avoid hitting the sedan.

There is something to this objection, insofar as Norma has a greater opportunity to avoid hitting the child than perhaps Pete does to avoid hitting the sedan. But making this distinction is not enough to show that tracing is necessary here. Recall that the purpose of tracing was to relocate the relevant control for some action to some prior action (or omission). Here, the objection means to show that Norma exercised control over her omission to check her mirrors, the sort of control Pete lacks over his omission to avoid the sedan (but which he did have over not getting drunk, say). If correct, then even if Norma exercised control over her omission to avoid hitting the child (as I have suggested), Pete will lack this control over hitting the sedan, and to preserve his responsibility for that action will require tracing back to some prior controlled omission.²²

Despite getting something seemingly right, there is also something off about this proposal. It seems to appeal to a dubious notion of control, one that is nothing at all like the control exercised in the cases of the soldiers, Rick, and Jill (or even Pete, on the recklessness interpretation). There, each agent manages, to some extent, the actions they take and has some mental bearing toward the outcome that results (either intending it or being aware but indifferent toward it). While there certainly are controlled omissions (e.g., declining to save a drowning swimmer), becoming distracted (Norma) or failing to consider certain possibilities (Pete on the negligence interpretation) aren't terribly plausible instances. As has been observed,²³ while we can attempt to locate control here, such a move would fly in the face of our experience. We all can too readily recall emerging from a distraction we had not realized we entered, or recognizing only after the fact our failure to notice some aspect of our situation or important feature worth considering. The assertion that all such failures are the result of control is thus implausible on its face, and would at least require significant and serious argumentation. Moreover, to sustain the significance of control here and to square such control with cases of intentional (and sub-intentional) action would be to admit that what matters is not the actual exercise of control, but rather the *opportunity* to exercise it.²⁴ But this amounts to a

²² Fischer and Tognazzini (2010) often appeal to this strategy of locating control in some past omission as a means of rebutting proposed difficulties for tracing. They claim that often it will be a controlled omission which grounds responsibility for some later harm.

²³ Sher 2006: 287-288. Sher argues that we cannot render consistent a control-based account with our strong intuitions of responsibility in cases that include failures of the sort discussed here. I take no stand on whether control-based accounts are problematic in general, holding instead to the weaker claim that tracing is a component of such accounts we should consider giving up.

²⁴ Holly Smith (1983) has argued that culpability for ignorance requires performing some act in the past that contributed to that ignorance. Such 'benighting acts', as she calls them, look to be references to tracing (in this respect Smith appears ahead of her time). One is responsible for the accident, say, because one neglected to have one's brakes checked on schedule, which led to the

radical departure from the notion of control originally sought by tracing backward in Pete's case. There, the idea was to find the same sort of control as that which typically grounds responsibility in some past action. This new proposal, however, replaces that typical control with a much different notion. It is not the opportunity to exercise control that grounds responsibility in Rick's case, or the soldiers'; what matters in those cases is that each in fact exercised control over their actions.

I conclude that drunk-driving cases need not rely on tracing to explain the agent's responsibility if interpreted on the model of negligence; indeed, tracing cannot do the requisite work. Explanations of negligence will not be able to locate a suitable past instance of control, or else they face an explanatory regress, or else they will force adoption of an altogether dubious notion of control, one not at work in ordinary cases of action. If this is right, then drunk-driving cases do not suggest a theoretical need for tracing. Either responsibility in drunk-driving cases can be explained by appealing to fairly pedestrian capacities of agents or else by a negligence mechanism. In neither case is tracing required.²⁵

5. Responsibility and control

What is required for an agent to control her action, in the sense relevant to ascriptions of responsibility? I have argued that it takes less than intentionally acting, for those who act knowingly or recklessly plausibly control their actions as well. But it takes more, it seems, than being negligent in one's action, as reflection on Norma's case revealed. I have claimed that if Pete's drunk driving is reckless, his hitting the sedan is sufficiently controlled, and thus there is no need to trace backwards at all. If Pete's drunk driving is negligent, neither the accident nor his continued drinking need be controlled in the relevant sense for him to be responsible. Thus, tracing would be of no help. On neither interpretation is tracing required.²⁶

But the significance of control extends far beyond its relevance to discussions of tracing. One finds control at the center of other discussions, from responsibility for personality traits (central to the dispute between voluntarists and attributionists²⁷), to moral luck,²⁸ to the traditional problem of free will captured in the debate over the ability to do otherwise.²⁹ There can be no doubt that control lies at the core of most theorizing about moral responsibility. There is much more to be said about control, then, than what can be adequately covered here. My aim has been to consider control in the context of a particular problem, the problem cases like drunk-driving pose, and to raise questions for the commonly accepted solution to that problem, tracing. Tracing has seemed so clearly necessary and so obviously correct for long enough that I cannot hope to convince its proponents outright by my argument here. However, I hope to have shown enough to raise serious doubts about the obviousness of tracing and its role in our

crash. As is hopefully obvious, such benighting acts have two interpretations, a reckless and negligent interpretation, just like Pete's case. And the results are the same.

²⁵ Of course, if we are out to square responsibility in cases of negligence with ordinary intentional action, we may be forced to give up control as a necessary condition. I take no stand on that issue here.

²⁶ Of course, even if there is a suitable tracing explanation for Pete's case, that alone would not show tracing to be *necessary* for explaining his responsibility. What proponents of tracing would need is to show that it is the only possible explanation.

²⁷ Voluntarists affirm whereas attributionists deny that control (usually in the form of choice) is necessary for all ascriptions of responsibility. (For examples of the former, see Levy 2005; McKenna 2008. For the latter, Scanlon 1998; Sher 2006; Smith 2008.) Proponents of tracing are thus apt to be considered voluntarists as well, though not all voluntarists need endorse tracing.

²⁸ Nagel 1976.

²⁹ For the case that free will requires the ability to do otherwise, see van Inwagen 1975. For the case against, see Frankfurt 1969. The ability to do otherwise may be thought of as one of the more contentious ways to characterize what controlling an action requires.

explanatory theories. By presenting a defensible account of a canonical type of tracing case that nonetheless does not appeal to tracing I undermined the claim that it is theoretically indispensable. Relying on standard resources from the philosophy of action, we can give an alternative explanation of drunk-driving cases, the prototypical case employed to call for tracing. At the very least, this fact undermines the supposition that tracing is the only available explanation and obviously correct. Even should tracing be ultimately vindicated, it should only occur after a more careful consideration of its virtues, not simply for lack of alternatives or intuitive plausibility.

If tracing is not necessary, this result is even more striking given tracing's ubiquity. Even its critics suppose it to be a necessary feature of our theories. But supposing its necessity is not to show that it is in fact indispensable. Undermining tracing's necessity also serves to sharpen the force of recent criticisms aimed at showing tracing to be a problematic component of theories of responsibility. It follows that if my argument is successful, and we can safely jettison tracing from our theories (at least with respect to the most standard of examples), these criticisms are no longer problems for everyone, but remain only for those theories committed to tracing.

In the final assessment, it may be that adequately explaining responsibility (or its lack) in certain cases requires some sort of tracing mechanism. For example, I have set aside discussion of so-called manipulation cases, which have garnered much attention of late as motivating examples, albeit of a different sort, for tracing's need. To fully abandon tracing, one would also need to take up these types of cases as well,³⁰ but settling the matter must await further inquiry.³¹ For now, I hope it clear that what the proponent of tracing cannot do is simply assume that tracing is required, nor especially make its inclusion a condition of adequacy on theories of responsibility. As significant, however, critics of tracing also cannot assume its necessity, for tracing's problems, whatever they might be, may not be everyone's problems after all.³²

References

- Bratman, M. (1987), *Intention, Plans, and Practical Reason*. Cambridge, MA: Harvard University Press.
- Ekstrom, L.W. (2000), *Free Will: A Philosophical Study*. Boulder, CO: Westview Press.
- Fischer, J.M. and M. Ravizza. (1998), *Responsibility and Control: A Theory of Moral Responsibility*. Cambridge, UK: Cambridge University Press.
- Fischer, J.M. and N. Tognazzini. (2009), 'The Truth about Tracing.' *Noûs* 43: 531-56.
- Fischer, J.M., R. Kane, D. Pereboom, and M. Vargas. (2007), *Four Views on Free Will*. Malden, MA: Blackwell Publishing.
- Frankfurt, H. (1969), 'Alternate Possibilities and Moral Responsibility.' *Journal of Philosophy* 66: 829-39.
- (2002), 'Reply to John Martin Fischer.' In S. Buss and L. Overton (eds) *The Contours of Agency: Essays on Themes from Harry Frankfurt*. Cambridge, MA: MIT Press.
- Haji, I. (2008), *Incompatibilism's Allure: Principal Arguments for Incompatibilism*. Peterborough, ONT: Broadview Press.
- King, M. (2009), 'The Problem with Negligence.' *Social Theory and Practice* 35: 577-595.

³⁰ For recent discussion of manipulation cases and their relevance specifically to tracing, see McKenna 2004, 2008; Mele 2008; and Vargas 2006.

³¹ This is my aim, indirectly, in 'Manipulating Responsibility' (ms). My argument implies that tracing is dispensable for manipulation cases because manipulated agents are still (directly) responsible for what they do.

³² I am indebted to John Martin Fischer, Mark Schroeder, Neal Tognazzini, Manuel Vargas, and a referee for the *European Journal of Philosophy* for their helpful comments.

- (ms), 'Manipulating Responsibility.'
- Levy, N. (2005), 'The Good, the Bad, and the Blameworthy.' *Journal of Ethics and Social Philosophy* 1: 1-16.
- McKenna, M. (2004), 'Responsibility and Globally Manipulated Agents.' *Philosophical Topics* 32: 169-92.
- (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*. New York: Oxford University Press.
- (2006), *Free Will and Luck*. Oxford: Oxford University Press.
- (2008), 'Moral Responsibility and History Revisited.' *Ethical Theory and Moral Practice* 12: 463-75.
- Nagel, T. (1976), 'Moral Luck.' *Proceedings of the Aristotelian Society* 50.
- O'Connor, T. (2000), *Persons and Causes: The Metaphysics of Free Will*. Oxford: Oxford University Press.
- Pereboom, D. (2001), *Living Without Free Will*. Cambridge, UK: Cambridge University Press.
- Scanlon, T. (1998), *What We Owe to Each Other*. Cambridge, MA: Harvard University Press.
- Sher, G. (2006), 'Out of Control.' *Ethics* 116(2): 285-301.
- (2009), *Who Knew? Responsibility Without Awareness*. New York: Oxford University Press.
- Smith, A. (2008), 'Control, Responsibility, and Moral Assessment.' *Philosophical Studies* 138: 367-92.
- Smith, H. (1983), 'Culpable Ignorance.' *Philosophical Review* 92: 543-71.
- Steinbock, B. (1985), 'Drunk Driving.' *Philosophy and Public Affairs* 14(3): 278-95.
- Strawson, G. (1994), 'The Impossibility of Moral Responsibility.' *Philosophical Studies* 75: 5-24.
- Tenenbaum, S. (forthcoming), 'Intention and Commitment.'
- Timpe, K. (forthcoming), 'Tackling Tracing's Troubles.'
- van Inwagen, P. (1975), 'The Incompatibility of Free Will and Determinism.' *Philosophical Studies* 27: 185-99.
- Vargas, M. (2005), 'The Trouble with Tracing.' *Midwest Studies in Philosophy* 29: 269-91.
- (2006), 'On the Importance of History for Responsible Agency.' *Philosophical Studies* 127: 351-82.
- Wallace, R.J. (1994), *Responsibility and the Moral Sentiments*. Cambridge, MA: Harvard University Press.
- Zimmerman, M. (1986), 'Negligence and Moral Responsibility.' *Noûs* 20: 199-218.