

TRACING AND THE EPISTEMIC CONDITION ON MORAL RESPONSIBILITY

Abstract: In “The Trouble with Tracing,” Manuel Vargas argues that tracing-based approaches to moral responsibility are considerably more problematic than previously acknowledged. Vargas argues that many initially plausible tracking-based cases of moral responsibility turn out to be ones in which the epistemic condition for moral responsibility is not satisfied, thus suggesting that contrary to initial appearances the agent isn’t morally responsible for the action in question. In the present paper, I outline two different strategies for responding to Vargas’ trouble with tracing. I then show how further consideration of the epistemic condition for moral responsibility renders tracing significantly less problematic than Vargas claims.

§ 1: INTRODUCTION

In “The Trouble with Tracing,” Manuel Vargas argues that tracing-based approaches to moral responsibility are considerably more problematic than previously acknowledged. Vargas argues that many initially plausible tracking-based cases of moral responsibility turn out to be ones in which the epistemic condition for moral responsibility is not satisfied, thus suggesting that contrary to initial appearances the agent isn’t morally responsible for the action in question. In the present paper, I outline two different strategies for responding to Vargas’ trouble with tracing; I argue that for one of these strategies to be successful, it must be grounded in the second strategy. I then show how further consideration of the epistemic condition for moral responsibility renders tracing significantly less problematic than Vargas claims.

§ 2: TRACING-BASED THEORIES OF MORAL RESPONSIBILITY

Despite their extensive divergence, most extant accounts of moral responsibility have two things in common. First, they hold that an agent must satisfy two distinct conditions to be morally responsible: a control condition and an epistemic condition.¹ The control condition is sometimes called the ‘freedom-relevant condition’ for moral responsibility. And free will is often taken to be the capacity or set of capacities needed to fulfill the control condition for moral responsibility. In other words, an agent acts freely when he controls his actions in the way needed

¹ There may well be other conditions required for moral responsibility as well, such as an authenticity condition or an ownership condition. Nothing in what follows hinges on the resolution of this issue.

for him to be morally responsible for that action. Despite receiving comparatively less attention, the satisfaction of the second condition is also widely thought to be necessary for moral responsibility. The intuition behind the epistemic condition² is that moral responsibility requires a certain kind of knowledge or belief about the nature of one's actions or choices and what could result from them. Peter van Inwagen captures the idea here as follows: "no one, I suppose, would seriously maintain that we can be blamed for *all* of the consequences of *any* of our acts.... Obviously, I can be blamed only for those consequences of my acts that are in some sense 'foreseeable'" (van Inwagen, 1989, 419). (Van Inwagen then goes on to say that there "is the dismally difficult question of what it is for a consequence of an act to be 'foreseeable' in the relevant sense" (421).)

The second, though admittedly less widespread, commonality among contemporary accounts of moral responsibility is they hold that an agent's moral responsibility for a particular action or choice can 'trace' back to previous times when the agent satisfies the relevant conditions for moral responsibility.³ Commitment to tracing is particularly prevalent among those approaches to moral responsibility which understand the control condition along the lines of sourcehood, rather than the mere having of alternative possibilities.⁴ The endorsement of tracing also cuts across the traditional incompatibilist/compatibilist divide. Here are a few representative samples:

Daniel Dennett:

"Here I stand," Luther said. "I can do no other." Luther claimed that he could do no other, that his conscience made it *impossible* for him to recant. He might, of course, have been wrong, or have been deliberately overstating the truth. But even if he was—perhaps especially if he was—his declaration is testimony to the fact that we simply do not exempt someone from blame or praise for an act because we think he could do no other. Whatever Luther was doing, he was not trying to duck responsibility (Dennett 1984, 13).

Robert Kane:

Incompatibilists need not deny Dennett's claim that Luther's "Here I stand" might be a morally responsible act, even if Luther "could have done no other" at the time of making it. To decide the issue, however, incompatibilists would insist that we know something about the circumstances and background of Luther's action that made him responsible or accountable for it. If Luther's affirmation

² The epistemic condition is also sometimes referred to as the knowledge condition. For reasons that will become clear later, I prefer to call this condition the epistemic condition.

³ Vargas defines tracing as follows: "tracing is the idea that responsibility for some outcome need not be anchored in the agent or agent's action at the moment immediately prior to [the] outcome, but rather at some suitable time prior to the moment of deliberation or action" (Vargas 2005a, 269).

⁴ According to Fischer and Ravizza, for example, it is the "tracing aspect of our theory of moral responsibility [which] implies that the theory is genuinely historical" (Fischer and Ravizza 1998, 196.). Chapter 7 of this work contains a very helpful and clear discussion of the difference between historical and time-slice approaches to moral responsibility.

did issue inevitably from his character and motives at the time it was made, then his moral accountability for it would depend on whether he was responsible *for being the sort of person he had become at that time*.... If we have no hesitation in saying that he was responsible for the final affirmation, I think it is because we believe that he was responsible through many past choices and action for making himself into the kind of man he was then.... If he is ultimately responsible for his present act, then at least some of these earlier choices or actions must have been such that he could have done otherwise with respect to them (Kane 1996, 39ff).

While not exhaustive, these two examples should be sufficient to show the widespread appeal to a tracing condition for various accounts of moral responsibility, regardless of where they stand on the compatibility issue. According to Manuel Vargas, “one of the nice features about tracing is that it is one of a few things to which nearly all parties in the debate about free will [and moral responsibility] appeal to with equal enthusiasm. Nothing about tracing presumes the truth of one or another view about the compatibility of moral responsibility on the one hand and determinism or the natural causal order on the other” (Vargas 2005a, 270).

§ 3: THE TROUBLE WITH TRACING

The widespread appeal of tracing-based approaches to moral responsibility, as well as the fact that it cuts across the typical incompatibilist/compatibilist divide, makes all the more pressing a recent challenge. In “The Trouble with Tracing,” Vargas argues that “the widely relied upon idea of ‘tracing’ in the theory of moral responsibility is considerably more problematic than has been previously acknowledged” (Vargas 2005a, 269).⁵ The idea of tracing could be applied to either of the two conditions whose satisfaction is necessary for moral responsibility mentioned above: the control condition and the epistemic condition. The appeals to tracing in the previous section are ones in which an agent’s control over a particular choice or action can be traced back to previous choices or actions over which the agent did have the requisite control. According to Robert Kane, for example, Luther was in control of his proclamation because the proclamation was the (purportedly) inevitable result of his having a particular character which, it is assumed, he had control over developing.⁶

In contrast, Vargas takes aim specifically at tracing as applied to the epistemic condition, or what he calls the knowledge condition (KC), for moral responsibility. He defines the epistemic condition as follows:

⁵ Similarly, Michael McKenna expresses a similar concern in a recent paper, where he writes that many tracing-based accounts “are not particularly realistic about the practical details of how choices earlier in life can have repercussions later in life” (McKenna 2008, 33). In Zimmerman 1997, Zimmerman argues that the conditions for culpable ignorance are satisfied far less frequently than is commonly supposed: “tracing culpable ignorance to such nonignorant origins would indeed require us to cut back fairly drastically on the frequency with which we ascribe culpability for ignorance” (418). It would be a worthwhile project to compare his account with Vargas’ worry over tracing; unfortunately, such a comparison cannot be made here.

⁶ See Kane 1996, 38-42, 77-78.

(KC) For an agent to be responsible for some outcome (whether an action or consequence) the outcome must be reasonably foreseeable for that agent at some suitable prior time (Vargas 2005a, 274).

Vargas intentionally formulates KC broadly, for his argument does not rely on a more fine-grained construal of the epistemic condition. He invites those “already committed to a more fine-grained account of the knowledge condition ... to substitute it for KC—it should not affect the outcome of the argument” (Vargas 2005a, 274) as “I suspect that similar difficulties can be raised for any plausible account of the knowledge condition” (Vargas 2005a, 284).⁷

Vargas proceeds by describing four cases; each case is described in two steps. The initial description aims at evoking the intuition that the agent is morally responsible for a particular action or choice, *x*; the second description elaborates the case in such a way as to present it as implausible that the agent could have reasonably foreseen the subsequent impact a previous choice or course of action would have upon *x*. While they differ considerably in their details,⁸ the underlying similarity is that “all [are] cases in which an agent has somewhat unexpectedly found him- or herself in a situation in which he or she may not count as a theory’s paradigmatic example of moral responsibility, but for which he or she is nonetheless seemingly responsible” (Vargas 2005a, 273). Though Vargas gives four cases, I’ll confine my attention in what follows to only two of those cases.⁹ These cases are as follows:

Jeff the Jerk:

Jeff is a middle aged middle manager in a mid-size company located somewhere in the Midwest. To him has fallen the task of alerting “downsized” employees of their new status as job seekers in a gloomy economy. That Jeff has the task is unfortunate for those about to be laid off, not only because they are about to lose their jobs, but—to add insult to injury—because Jeff is jerk. He is rude and inconsiderate about the feelings of others. And, he is unreflective about it. When people react poorly to his behavior (something they avoid doing because he is large, imposing, generally unsympathetic, and even a little frightening) he always writes it off as a shortcoming on the part of others. One afternoon, his superiors tell him that he needs to give notice to a group of long-time employees

⁷ According to both Vargas and Bob Kane (in personal correspondence), it is likely that the epistemic condition will itself be neutral to the compatibilism/incompatibilism debate.

⁸ The primary difference among the four cases is their respective focus on character traits (Jeff the Jerk), dispositions that are less robust than character traits (Britney), the interaction among traits (Paulina), and the consequences of interaction among traits (Ruben). Despite these differences, I agree with Vargas when he says “the success of my argument does not depend on your agreement with my treatment of each and every one of these case. Depending on your account of moral responsibility (on the assumption that it relies on tracing), perhaps one or two of these cases can be rescued. All my argument requires is that at least one case (though I suspect that for most theories, it will be most of the cases) raises difficulties for your account” (Vargas 2005a, 282.).

⁹ For considerations of space, and for reasons enumerated below in footnote 25, I do not consider Vargas’ cases of Britney the Bride and Ruben the Unfortunate here.

that they will be laid off. He does tell them, but in an altogether rude, and insensitive fashion... (Vargas 2005a, 271).

As it is for many of us, high school was a crucial formative period in Jeff's life. He does not look back at it with any special fondness. In fact, he has largely forgotten the details. Nonetheless, it played an important role in creating the person who would become Jeff the Middle Management Jerk.

When Jeff was 15, he realized that he was having much less success with members of the opposite sex than he wanted to have. Over time, and through the usual fallible mechanisms of belief acquisition, he came to believe that the only males who consistently had success at gaining the attention of female classmates were those at least *we* might describe as jerks. In Jeff's hormone-ridden 15-year-old mind, this putative insight, coupled with a somewhat enterprising disposition that later served him modestly well in the business world, led him to adopt a plan for self-improvement. To the extent to which he was able, he inculcated in himself all the behaviors and attitudes that we would perceive to be jerk-like, and therefore, ultimately conducive to success with his female classmates. Now this initial decision to undertake the program of self-improvement was not an obvious choice for Jeff. He was worried about what his friends would think when he started attempting to behave differently. Would they make fun of him, noting that he was pretending to be something he wasn't? What if people found out why he was doing it? Could he overcome the shame and humiliation that would result from his female classmates learning of his subterfuge? These were the sorts of thoughts that Jeff had while he deliberated. After a period of uncertainty, however, Jeff decided to undertake the plan. And so he set about becoming, if not a jerk, at least jerk-like.

With surprisingly little effort, he succeeded. In fact he more than succeeded—it didn't even take the whole academic year for him to go from being jerk-like to being a full-on jerk. Part of the reason was undoubtedly rooted in his social context. The context was such that there was little cost to behaving in jerk-like ways when he would have been sensitive to those sorts of feedback. (He was a "latch key" kid, largely ignored by his permissive parents who chalked up his increasingly rude behavior to being a typical teenager. So, they tuned it out and treated him in a largely hands-off fashion). Moreover, what feedback he did receive and did care about was overwhelmingly positive. For whatever reason (perhaps it was merely the growth spurt that struck around the same time), he came to have tremendous success at attracting the attention of the opposite sex. However, at no point during the process of becoming a jerk, and certainly at no point before he undertook the program of

becoming jerk-like, did Jeff even conceive that his plan for personal improvement would include in its outcomes that he would some day lay off employees in a despicable fashion. But it did, even though this was not at all reasonably foreseeable given his age and context.

Of course, these are all plausibly contingent facts about Jeff's history. Had he been some place else, perhaps as close as the school on the other side of town, things would have worked out differently. The people and social context would have been such that he would have ultimately abandoned the plan to become a jerk. Perhaps he would have been found out. Perhaps someone would have convinced him that this would have other deleterious effects on his life or on others. Perhaps someone might have even convinced him that it would lead to him becoming a quietly disliked middle-aged middle manager at a mid-sized company in the Midwest. But, alas, he was in the wrong place at the right time. Though he did become a jerk, and though he did become a jerk freely, and though he did foresee of his becoming a jerk that it would help him with his aspirations regarding the opposite sex, given his history and context it was not reasonably foreseeable to his hormonefogged 15-year-old mind that undertaking this course of action would lead to his becoming the quietly disliked sort of person who would terminate the employment of people in unnecessarily rude and degrading ways (Vargas 2005a, 275f).

Paulina the Paralyzed:

Paulina is visiting Florida for the first time. She takes her 2-year-old son, Paul, on a walk in a park just outside of Tallahassee. They eventually find a small clearing and decide to sit down and eat lunch. After a while, Paul gets up and starts wandering around with half-eaten food in his hands, on his face, and on his clothes. Paulina, being a responsible mother, keeps an eye on Paul and warns him against straying too far. Near the edge of the clearing, about 20 feet away, Paul bends over to pick up a rock that catches his attention. Paulina watches him do this, but at that moment realizes that there is an alligator staring at Paul from about 30 feet away. Paulina is paralyzed by surprise and some degree of fear. She feels a rising sense of panic, but remains frozen as the alligator start to move in the direction of Paul. Alas, the alligator moves surprisingly quickly, and snaps up poor Paul. Only then does Paulina finally unfreeze, and then she screams... (Vargas 2005a, 272).

Paulina has a standing habit of taking Paul for walks in a park in the afternoon. It was the first time she had been to Tallahassee, and because she just never bothered to think about it, she did not have any expectations about the native fauna. When she did see the

alligator, she was paralyzed because of several interacting factors. First, she did not expect to see an alligator. She was simply surprised. Second, it was big—a lot bigger than she thought alligators were. Third, it brought forth a visceral fear she hadn't felt in years—a fear of live reptiles with lots of big teeth. The absence of any one of these conditions would have meant that she moved to attempt to save Paul, but the confluence of these conditions is what left her paralyzed, swamping her desire to try and save her son for just long enough. Here's the thing: she could have undergone therapy that would have reduced or mitigated her fear of reptiles with lots of teeth. After all, her father was a therapist and her mother a herpetologist—so, she grew up knowing that she was afraid of these creatures, and she knew that if she really wanted to she had available to her the means to overcome her fear. But, the last time she thought about this fear she had no reason to believe that it would affect her ability to save her son (indeed, at the time, she didn't even have plans to have children). So, she decided to let sleeping fears lie (Vargas 2005a, 281).

All of Vargas' cases involve “an action (or omission) that is brought about in large part by the presence of a disposition, character trait, habit or other non-deliberative aspect of the agents” (Vargas 2005a, 275). Vargas focuses on non-deliberative cases because there are going to many cases where “the non-deliberative source of behavior was acquired or retained under conditions where the agent could not have reasonably foreseen the later consequences of having that disposition, habit, or character trait” (Vargas 2005a, 275). In these cases, it looks as if the agent fails to satisfy standard account of responsibility despite the intuition that the agents in these cases are responsible. I think that Vargas is correct in thinking that it is non-deliberative cases such as these examples which are particularly problematic for tracing, as it is much easier to reconcile our intuitions and our theory about moral responsibility in explicitly deliberative cases. Consider a simple case in which an agent fails to satisfy the epistemic condition for moral responsibility at a particular time, but where that failing can be traced to a prior deliberative act by the agent. Such would be the case, for example, for the student who fails to know material covered on a test in class because the previous night he had deliberated about whether to study the course material or go to the frat party and decided to pursue the latter course of action. There does not seem to be any conflict between a theory of moral responsibility involving a tracing condition and our intuition about the agent's moral responsibility in this sort of case. In order for an agent to be morally responsible in cases where his responsibility traces back to an earlier time, “we have to find a prior moment when he could both act freely [thereby satisfying the control condition] *and* reasonably foresee the outcome [thereby satisfying the epistemic condition]” (Vargas 2005a, 277). But in non-deliberative cases those moments at which the agent satisfies the control condition appear to be such that the agent

couldn't reasonably foresee the later impact, thereby violating the epistemic condition on moral responsibility.

Vargas claims this 'trouble with tracing' will infect, to some degree or other, all accounts of moral responsibility that admit of tracing:

if restrictive accounts of free will are true, then given the extent to which restrictive accounts must rely on tracing (and in particular, tracing back to the formation of character traits and habits), the difficulties I have raised for satisfying the epistemic condition strike especially deep. Even on permissive views [according to which we are often free], though, the difficulties I raise do not disappear. . . . It appears that even the most permissive views that countenance tracing will have to admit that in a fair number of tracing-invoking cases the knowledge condition on moral responsibility will not be satisfied (Vargas 2005a, 286).

According to Vargas, the more theoretically attractive one's combined account of tracing and the knowledge condition are, the less frequently agents will satisfy the conditions for moral responsibility. He also describes the problem as a "difficulty of balancing, on the one hand, a plausible account of tracing and the knowledge condition with, on the other hand, the full scope of cases we think of as intuitively responsible" (Vargas 2005a, 284).

§ 4: AVENUES OF RESPONSE

How should one respond to Vargas' provocative argument? There are a number of available lines of response—some of which seek to accommodate the 'trouble' Vargas raises and some of which seek to minimize it. One could, for example, jettison the possibility of tracing from one's account of moral responsibility. This strategy, however, seems unduly drastic; it is hard to see, for example, how one could account for a drunk driver's being responsible for running over a pedestrian without a tracing clause. The drunk driver's level of inebriation could be high enough that he literally could not control hitting the pedestrian at that time; but surely it is a mark against a theory of moral responsibility that it fails to hold drunk drivers responsible. Alternatively, one could accept Vargas' critique of tracing as applied to moral responsibility and substantially restrict the range of cases in which we hold agents morally responsible. We might think, for example, that while tracing can occur, the range of cases in which the relevant conditions on tracing are satisfied is relatively narrow. Unlike the first option, this approach would hold that the drunk driver was morally responsible, so long as he clearly satisfies both the epistemic and the control conditions for responsibility at the time that he became inebriated. But like the first option, this approach would still appear to vitiate our moral intuitions in a large number of cases, including the four that Vargas delineates. While Vargas might not be disinclined about such a wholesale revision of our intuitions about moral responsibility,¹⁰ the severity of such revision should, I think, give us pause.

¹⁰ See Vargas 2004 and 2005b for a discussion of revisionism.

One way of showing that such revision is not needed would be by meeting the trouble head-on, as it were, and showing that an account of tracing as applied to moral responsibility can be given which avoids Vargas' critique. Here, at least two general strategies are available.¹¹ According to the first strategy, which we might call the *Denial Strategy*, one would show that each of the purported cases are, despite initial appearances, not counterexamples to the plausibility of tracing. That is, one would deny that the cases presented by Vargas really are problematic for tracing-based accounts of moral responsibility. John Fischer and Neal Tognazzini pursue the Denial Strategy in a recent paper, "The Truth about Tracing" (Fischer and Tognazzini 2009). There is much in Fischer and Tognazzini's discussion of Vargas' four cases with which I am sympathetic. In fact, as I show in § 7 below, I agree with many of their claims when it comes to evaluating the individual cases that Vargas mentions. As a result, I think that they have successfully engaged in some advantageous denial. But I think there the second strategy, which I pursue below, can actually bolster and undergird their Denial Strategy.

I refer to this other strategy as the *Bring It On Strategy*, insofar as it directly responds to the following challenge posed by Vargas:

a fourth complaint is that what these cases show is that KC is simply an inaccurate construal of the knowledge condition. To this complaint, my reply is George W. Bush's (and perhaps as ill-advised): Bring it on. That is, I suspect that similar difficulties can be raised for any plausible account of the knowledge condition (Vargas 2005a, 284).

This strategy seeks to refine the epistemic condition in such a way that the purported troublesome cases (or at least a significant sub-set of those cases) will fall *en masse*. In 'bringing it on', this second approach thus aims to nullify Vargas' central claim that tracing-based approaches face widespread problems in simultaneous satisfaction of the control and epistemic conditions on moral responsibility. Vargas himself indicates refining the epistemic condition seems to be more promising than abandoning tracing altogether. But such a task, he thinks, will not be easy.

I agree with Vargas that the Bring It On Strategy will be harder to engage in than the Denial Strategy, but I also think it would be worthwhile if it could provide a general schema by which the purported trouble can be eliminated. The more systematic nature of 'bringing it on' suggests a certain desirability that I want to pursue in what follows—though later I shall also argue that the Denial Strategy, if it is to be ultimately successful, will be grounded in the Bring It On Strategy.

§ 5: EXTANT EPISTEMIC CONDITIONS

As noted earlier, Vargas' challenge to tracing uses a fairly thin epistemic

¹¹ There may be other strategies as well. Furthermore, it may also be the case that the two strategies considered here can complement each other, as I suggest below.

condition on moral responsibility:

(KC) For an agent to be responsible for some outcome (whether an action or consequence) the outcome must be reasonably foreseeable for that agent at some suitable prior time (Vargas 2005a, 274).

While a virtue of Vargas' use of KC in discussing tracing is that such a thin condition doesn't require substantial commitments, its very thinness is also problematic. So in this section I want to canvas other extant accounts of the epistemic condition on moral responsibility as a way of giving more positive content to the epistemic condition. After doing this, I will try to show how a revised epistemic condition gives plausibility to the Bring it On Strategy.

In a recent book, Ish Haji considers the following epistemic condition on moral responsibility:

E-CONDITION 1: Agent, *S*, is morally blameworthy for action, *A*, only if *S* knows that *A* is morally wrong; and *S* is praiseworthy for *S* only if *S* knows that *A* is morally obligatory or permissible (Haji 2008, 89).

Haji quickly rejects such a condition as too strong: "Surely, you can be blameworthy for an action even though you did not *know* that it was wrong, had good reason, though, to believe that it was wrong, and intentionally performed the act despite your belief that you would be doing intentional wrong in performing it" (Haji 2008, 90). He then suggests the following, more plausible, candidate for the epistemic condition:

E-CONDITION 2: Agent, *S*, is morally blameworthy for action, *A*, only if *S* does *A* at least partly on the basis of the belief that *A* is morally wrong; and *S* is praiseworthy for *A* only if *S* does *A* at least partly on the belief that *A* is morally obligatory or permissible (Haji 2008, 90).

When coupled with the uncontroversial premise that *s* is morally blameworthy for *a* only if *a* is morally wrong and *s* is morally praiseworthy for *a* only if *a* is morally obligatory or permissible, Haji suggests that E-CONDITION 2 is a more plausible principle. He thinks that the epistemic requirement for moral responsibility must be at least as strong as E-CONDITION 2 (even if it needn't be as strong as E-CONDITION 1). "I know of no discussion on the epistemic requirements of moral responsibility in which it is maintained that *no* sort of belief or awareness constituent is any part of the epistemic constraint" (Haji 2008, 99).

Like Haji, I think there is good reason not to require actual knowledge in the epistemic condition.¹² Nevertheless, I think that Haji's E-CONDITION 2 is still too strong. On Haji's E-CONDITION 2, an agent could fail to be morally responsible for a particular heinous action by simply failing to believe that what she's done is morally wrong, even if she *should have* had such a belief. That is, Haji's condition fails to take account of culpable ignorance—those cases in which an

¹² Similarly, Carl Ginet writes that requiring knowledge would be 'requiring too much'; see Ginet 2000, 269f.

agent is ignorant in the relative way, but where her ignorance is blameworthy and supports, rather than undermines, her moral responsibility.¹³ While differentiating culpable ignorance from non-culpable ignorance is not always an easy task, the epistemic requirement needs to be sensitive to the difference.¹⁴ We shall return to the relationship between culpable ignorance and tracing in section 7 below. But for now I shall assume that an account of the distinction between culpable and non-culpable ignorance can be provided.

Unlike Haji, Michel Zimmerman makes an allowance for non-occurrent beliefs, but it seems to me that his allowance is not actually strong enough. Zimmerman writes:

while beliefs presumably can be merely dispositional, it is, I believe, occurrent beliefs about wrongdoing that are, with one possible exception, required for culpability; and such beliefs involve advertence.... The one possible exception is this: it may be that routine or habitual actions are performed for reasons to which one does not advert. It may also be that some people engage in deliberate wrongdoing in a routine or habitual, and hence inadvertent, manner (Zimmerman 1997, 421f).

Zimmerman is right that moral habits or dispositions can lead one to fail to attend to beliefs about the moral status of one's actions. An example might be young teenagers who spray paint a neighborhood fence without taking the time to consider that doing so will cost the residents money to repair it. But sometimes we are inclined, and I think rightfully so, to hold an individual responsible for not considering the moral status of an action even if that failing isn't the result of a flawed habit. Think of instances where an individual acts 'out of character'. Consider, for example, a normally attentive and truthful husband who, perhaps because he is rushed or distracted, tells his wife something that he knows is false without thinking about the wrongfulness of what he's doing. The husband knows that lying to one's spouse is wrong, but on this occasion, because of being hurried or distracted, simply fails to consider 'telling my wife this is an instance of lying, and therefore wrong'. Of course, if he were to stop and think about it, he'd have the relevant occurrent belief about the moral status of his action. Short of there being some other excusing factor for his behavior, we would not declare him non-responsible for his lie simply because he is acting contrary to his normal moral character. So the range of cases in which occurrent beliefs are not required for moral responsibility are wider than suggested by Zimmerman.¹⁵

Carl Ginet's "The Epistemic Requirements for Moral Responsibility" contains an extensive discussion of the epistemic condition. Ginet's proposed epistemic condition is complex. First some terminological explanation is needed. In his discussion, Ginet uses 'know*' and cognates as follows: to say that an agent

¹³ The literature on culpable ignorance is, not surprisingly, both vast and diverse in its conclusions. For two particularly helpful papers on the topic, see Smith 1983 and Zimmerman 1997.

¹⁴ In contrast to what I suggest above, W. D. Ross suggests that ignorance, even if culpable, fully excuses an agent from moral blame. See Ross 1939, 157-64.

¹⁵ For a treatment of culpable ignorance as it relates to occurrent vs. implicit belief, see Guerrero 2007.

knows* some truth, T , at a particular time implies that the agent actively (that is, occurrently) and truthfully believes T at that time, and has proper justification for this belief (Ginet 2000, 270). With this in mind, here then is Ginet's epistemic condition, the satisfaction of which he takes to be necessary for moral responsibility:

- (K) either (i) S knew* at t_1 that his moving in way M then would or might bring about (or, the case of omission, prevent) a harm of sort H , or
- (i) S did not then know* this but there is a sequence of one or more prior acts (or omissions) that ends with the act (or omission) at t_1 and is such that
- (a) each member before this last member benights the subsequent member,
 - (b) the first (earliest) member of the sequence was *not* a benighted act (or omission)—at the time S knew* that her act (or omission) would or might lead to the sort of harm it in fact let to, namely, the benighted act (or omission) that is the next member of the sequence—and
 - (c) at the time of each benighted act (or omission) in the sequence S should have known* (was blameworthy for not knowing*) that it would or might lead to the sort of benighted act (or omission) that it in fact let to in the next member of the sequence (Ginet 2000, 275).

There are number of important things to note about Ginet's (K). First, disjunct (i) doesn't required that S know* that M *would in fact* lead to H , just that it 'would or might'. Like Ginet, I think that requiring an agent to understand what 'would' happen sets the epistemic bar too high. Suppose that an individual is contemplating throwing a rock off the top of the Empire State Building during morning rush-hour. While one cannot be certain that the rock would strike a passerby far below, surely the possibility of 'getting lucky' and the rock landing safely in the street does not render the individual morally blameless for the toss.¹⁶ Furthermore, slightly modifying the case so that the rock was thrown at 8:17 pm, when perhaps the chance of hitting a pedestrian below is slightly less than 50%, would not change our attribution of blameworthiness. For this reason, it is be better to build the epistemic condition around what 'might' happen.

On the other hand, an agent's being aware of the mere possibility of an outcome sets the epistemic bar too low, for there are countless known possible outcomes of nearly every action, and it would be inappropriate to hold agents responsible for all of them. So, just how likely does a result need to be in order to

¹⁶ I originally owe this point, as well as the example, to Alicia Finch. I have subsequently found the same point made using a similar example in Ginet 2000, 269f.

be relevant to the epistemic condition? As Fischer and Ravizza note, the ‘might’ will need to be context sensitive:

In some contexts, it seems appropriate to hold an agent responsible for a later action (or omission or consequence) that is extremely unlikely to occur, whereas in other contexts the extreme unlikelihood of (say) the action seems to rule out responsibility. This makes it reasonable to think that a full and explicit tracing approach would not simply specify a degree of likelihood that is always employed straightforwardly to ascertain responsibility; rather, the degree of likelihood employed by the tracing approach would need to be context-relative (Fischer and Ravizza 1998, 50 note 21; similarly Ginet says that the possibility must be “non-negligible” (2000, 270)).

(More about the ‘might’ in the following section.)

Second, disjunct (ii) of (K) adopts ‘benighting’ language from the work of Holly Smith. In her insightful discussion of culpable ignorance, Smith notes that such cases

all involve a sequence of acts: an initial act, in which the agent fails to improve (or positively impairs) his cognitive position, and a subsequent act in which he does wrong because of his resulting ignorance. For convenience I shall call the initial act the ‘benighting act’ and the subsequent act the ‘unwitting wrongful act’ (Smith 1983, 547).

Ginet refers to the ‘unwitting wrongful act’ as the ‘benighted’ act. Such language easily maps onto the tracing cases under consideration, in which the agent’s responsibility for a particular *benighted* action is traced back to a previous *benighting* act. However, like Haji’s account, both (i) and (ii) of Ginet’s (K) requires the agent to have an occurrent belief about the moral status of her action, which sets the epistemic bar too high.¹⁷ Furthermore, it’s not clear that it avoids all fo the problems raised against the other epistemic conditions canvassed here.

§ 6: REVISING THE EPISTEMIC CONDITION

In the previous section, I canvassed a number of extant epistemic conditions for moral responsibility, focusing on relevant strengths and weaknesses of those approaches. In the present section, I want to use the results from that discussion to develop what I think is a plausible epistemic condition on moral responsibility with more positive content than Vargas’ KC. In the following section, I will explore how this epistemic condition can respond to the troubles for tracing Vargas raises via a version of the Bring it On Strategy.

I take the following to be desiderata for constructing a satisfactory epistemic condition for moral responsibility suggested by the previous discussion:

- (1) the condition should require not knowledge, but mere reasonable

¹⁷ While part (c) of (ii) doesn’t require occurrent belief, it is required in part (b) as well as in (i).

belief, both about the moral status of the benighted action (or actions), as well as regarding the relationship between the benighting action (or actions) and the benighted action,¹⁸

- (2) in order to account for culpable ignorance, the condition should not require occurrent belief regarding the moral status of the benighting action (or actions), the benighted action, or the relationship between the benighting action (or actions) and the benighted action,
- (3) the condition should not require that the relevant beliefs regarding the benighting action (or actions) involve what ‘will’ or ‘would’ follow from such an action, but rather need only be about what ‘might’ happen as a result of the benighting action (or actions), and
- (4) the condition should allow for the ‘might’ in the previous desideratum to vary according to the expected value of the act in question; sometimes the probability of our actions leading to something bad needn’t be very high for us to still be blameworthy for those actions given the seriousness of the outcomes, but other times the bar will be higher.

This fourth desideratum is unfortunately vague; but the intuition driving it is, I think, clear enough. According to (4), the sense of ‘might’ involved in ‘y might cause x’ is best not thought of simply as a function of the probability of y’s occurring leading to the occurrence of x, but instead as a function of both the probability and importance of x’s occurring given y’s occurrence. Here it may be helpful to consider some remarks by Joel Feinberg in the context of legislation:

It will not do to say simply that only harm whose occurrence as a consequence of a given activity is *more probably than not* will justify preventive coercion. If the harm in question is very great, then a very small likelihood of its occurrence will do.... Conversely, if the harm in question is relatively small ... it may be worth running a substantial risk of its occurrence to avoid using coercion.... The inverse variation between magnitude and probability of harm can be expressed in concise formulas, the greater the probability of harm, the less grave the harm need be to justify coercion; the greater the gravity of the envisioned harm, the less probable it need be (Feinberg 1987, 190f).

To illustrate this idea, compare two actions: *A* and *B*. Suppose, on the one hand, that the performance of action *A* has a relatively low probability of leading to a bad, but fairly minor, consequence. For instance, perhaps Alvin is juggling

¹⁸ Whether or not these beliefs must be true will depend on other aspects of one’s moral theory. Aquinas, for example, held the controversial view that an erring conscience binds, such that one’s beliefs about the moral wrongness of an action need not be true in order for the agent to be morally blameworthy about doing that action (see *Summa Theologiae* IIaIIae 19.5). Those who find such a claim to be problematic are welcome to make the further requirement that the beliefs involved here are true.

expensive watches he has borrowed from passers-by in a busy train station without telling them what he would be doing with them. While he's an accomplished juggler, he sometimes (but not always) drops what he is juggling. Thus, there is a chance that he could drop and damage others' property. Suppose, on the other hand, that the performance of action *B* has the same probability of leading to a much more serious outcome. Perhaps Bertha is an equally accomplished juggler as Alvin is, and thus has the same chance of dropping one of her props. But rather than juggling watches borrowed from passersby, Bertha is juggling glass vials of mustard gas. Given that the consequences of Bertha's unlikely drop are significantly more severe than the consequences of Alvin's equally unlikely drop, it is very plausible that Bertha should be required to pay more attention to the possible, even if unlikely, outcomes of her juggling in the train station. Or, to use an example involving character formation, we should be more sensitive to patterns of action that might lead us to become abusive of our children than to patterns of action that might lead us to become insensitive to publication deadlines given the greater moral weight in the former case.¹⁹

Before attempting to formulate an epistemic requirement for moral responsibility based on these desiderata, one further desideratum should be added to the others. In his discussion of ignorance, Thomas Aquinas notes that there are two kinds of ignorance that do not *per se* undermine an agent's moral responsibility. The first of these, which he calls 'consequent ignorance', involves cases where an agent is ignorant of a morally relevant fact where that agent is responsible for said ignorance, such as when a student is blameworthy for his performance on an exam given that he failed to study for it.²⁰ In this sort of case, as mentioned earlier, the responsibility for one's ignorance can be traced back to a previous action which led to and resulted in the later ignorance. In addition, Aquinas thinks that there is another species of ignorance, called 'concomitant ignorance', which also

¹⁹ Something similar is reflected in the law. One's degree of negligence, understood as involving objective fault, is directly proportional to the gravity of the harm that foreseeably would result from one's conduct and is inversely proportional to the probability of such harm occurring. While civil negligence involves a deviation from the reasonable standard of care, criminal negligence involves a gross deviation from that standard. Furthermore, according to tort law, "a person acts 'recklessly' if he takes a very substantial and unjustifiable risk. So understood, 'civil negligence', 'criminal negligence', and 'recklessness' lie on a continuum: each involves unjustifiable risk-taking; they differ only in respect to the degree of the actor's deviation from the standard of due care" (Dressler 2001, 133.). A referee points out that Dressler's interpretation of these concepts is "somewhat non-standard, since the more standard interpretation views 'negligence' as involving non-advertence to the risk, while 'recklessness' involves conscious advertence to the risk."

²⁰ To be more exact, Aquinas differentiates two kinds of consequent ignorance: "[Such] ignorance is 'consequent' to the act of the will, in so far as ignorance itself is voluntary: and this happens in two ways.... First, because the act of the will is brought to bear on the ignorance: as when a man wishes not to know, that he may have an excuse for sin, or that he may not be withheld from sin.... And this is called 'affected ignorance'. Secondly, ignorance is said to be voluntary, when it regards that which one can and ought to know: for in this sense 'not to act' and 'not to will' are said to be voluntary.... And ignorance of this kind happens, either when one does not actually consider what one can and ought to consider; this is called 'ignorance of evil choice', and arises from some passion or habit: or when one does not take the trouble to acquire the knowledge which one ought to have; in which sense, ignorance of the general principles of laws, which one ought to know, is voluntary, as being due to negligence" (*Summa Theologiae* I-II.2.6.8).

does not undermine responsibility. Concomitant ignorance occurs when an agent is, through no fault of his own, ignorant of some morally relevant feature of an action but would have acted in the same way as he actually did even had he not been ignorant. Aquinas gives the following sort of example.²¹ Say that Tamler wishes to kill his arch-enemy, Dan, and plans on making an assassination attempt on the latter's life in the near future using a bow and arrow. To prepare, Tamler heads to Muir Woods to hone his archery skills. Tamler sees something in the underbrush, but cannot make out what it is. Thinking it is a deer, he takes careful aim and lets an arrow fly. The arrow flies true and hits its mark. Hearing a squeal from his impaled victim, Tamler rushes over to find that his victim was none other than Dan who, unbeknownst to Tamler, was hiking in the area. Tamler's shot proves lethal and much to his pleasure, Dan is dead. In this case, Tamler's ignorance of what is rustling in the underbrush does not excuse his killing of Dan because *had Tamler known* that it was Dan, and not a deer, in the brush he still would have taken the shot.²² We might take Aquinas' insight and use it to generate a fifth desideratum:

- (5) the condition should allow for an agent's being morally responsible in cases involving concomitant ignorance.

Guided by these constraints, I'd like to suggest the following epistemic condition for moral responsibility as an improvement over Vargas' KC—where various aspects of this condition are understood in light of the above comments. (Unfortunately, this improved epistemic condition, which I shall call IEC, is even more unwieldy than Ginet's (K); while I am unhappy with this complexity, I see no way around it given the issues involved.) Here, then, is the proposed condition:

- IEC: For an agent to be morally responsible for an action or the consequence of an action x at time t :
- (i) the agent must at t have reasonable belief about the moral status of his x -ing, or be culpably ignorant at t of the moral status of his x -ing; or
 - (ii) there must be some previous time t_1 at which the agent
 - (a) had reasonable though again not necessarily occurrent belief about the moral status of his x -ing or was culpably ignorant of the moral status of his x -ing;
 - (b) a reasonable though not necessarily occurrent belief that his y -ing might lead to his x -ing in the future (where the 'might' is understood as a function of expected value), or was culpably ignorant of

²¹ See *Summa Theologiae* IaIIae 6.8. Alexander Guerrero also thinks that cases of this sort are ones in which the agent is morally responsible; he refers to them as involving 'actions from callous indifference'. See Guerrero 2007.

²² Presumably, a parallel case could be given involving cases of concomitant ignorance of morally good actions as well.

- the fact that his *y*-ing might lead to his *x*-ing in the future;
- (c) freely *y*-ed; and
 - (d) the agent's freely *y*-ing at *t*₁ leads to his *x*-ing at *t*;
- or
- (iii) the agent is non-culpably ignorant of either
 - (a) the moral status of his *x*-ing,
 - (b) the moral status of his *y*-ing, or
 - (c) the fact that his *y*-ing might (again, understood as a function of expected value) lead to his *x*-ing in the future,
- but even if the agent had been aware of (a), (b), and (c), he still would have performed *y*.²³

I think that something along the lines of IEC plausibly captures the epistemic condition on moral responsibility. In the following section, I want to explore a number of implications IEC has for Vargas' challenge for tracing-based accounts of moral responsibility.

§ 7: TRACING REVISITED

Remember from § 3 that at the heart of Vargas' challenge for tracing-based accounts of moral responsibility is the claim that many initially plausible cases of tracing turn out, upon further investigation, to be cases in which there is no suitable prior time at which the agent satisfies both the control and epistemic conditions for moral responsibility. If there is no such prior time to which the agent's responsibility can be traced, then the agent in question will not be morally responsible for the later action, contrary to the initial intuition. If tracing-based approaches "have to admit that in a fair number of tracing-invoking cases the knowledge condition on moral responsibility will not be satisfied" (Vargas 2005a, 286), tracing will not play the widespread role in moral responsibility that many theorists think that it plays. While I think that there are cases where it will turn out that an agent's moral responsibility cannot be traced back to an earlier action or set of actions given that the agent will not satisfy both the control and epistemic conditions at this previous time, I do not think that the class of such cases is as wide as Vargas suggests. For this reason, I think that tracing faces far less trouble than does Vargas.

According to Vargas, there is an inherent tension between the epistemic condition and tracing-based accounts of responsibility:

the more theoretically attractive we make our combined accounts of

²³ I thank Alicia Finch especially for helping me formulate IEC in this way. I also intend IEC to apply only to morally significant actions, those actions for which the agent is either morally praiseworthy or morally blameworthy. Morally neutral actions are actions for which, in principle, the agent could be morally responsible even if she fails to satisfy (i), (ii), and (iii). The reason for this limitation is that modifying IEC to apply to morally neutral actions would make it even more unwieldy. I thank Al Mele and Matt Talbert for raising this issue.

tracing and the knowledge condition, the less frequently we seem to have free actions (and/or moral responsibility). Conversely, the broader range of cases we seem to insist are free (and/or responsible), the less attractive (and perhaps *ad hoc*) our account of tracing and the knowledge condition seems to become.... What this points to, I think, is the difficulty of balancing, on the one hand, a plausible account of tracing and the knowledge condition with, on the other hand, the full scope of cases we think of as intuitively responsible (Vargas 2005a, 271 and 284).

Is IEC *ad hoc*? It doesn't seem to me to be so; rather, the various components of it are motivated by the desiderata outlined in the previous section.²⁴ In the remainder of this section, I first draw out some general implications that IEC has for tracing cases. I then show how IEC can be used to justify Fischer and Tognazzini's responses to Vargas' purported troublesome cases for traces.

First, IEC maintains the agent's moral responsibility in those cases where the agent is ignorant but culpably so. Of course, spelling out those conditions under which ignorance is culpable is not easy. There are a number of competing accounts available in the literature, and it isn't clear that any of them has succeeded in providing a complete set of necessary and sufficient conditions for culpable ignorance.²⁵ But any plausible account of moral responsibility faces the same lacuna, so the fact that IEC doesn't spell out such conditions ought not count too much against the present proposal. This, however, won't suffice to address all the worries about the relationship culpable ignorance and IEC, for there remains a serious concern about circularity. In section 5 above, I assumed that an account of the distinction between culpable and non-culpable ignorance could be provided. But if one's account of culpable as opposed to non-culpable ignorance itself appeals to tracing, then the overall account of the epistemic condition on moral responsibility will be circular. This is a serious worry, and one that I'm not fully able to put to rest since, as mentioned before, I am not here giving an account of the nature of culpable ignorance. But as a partial defense, let me say the following.²⁶ It is not obvious to me that distinguishing between culpable and non-culpable ignorance requires an appeal to tracing. It is at least plausible that a counterfactual account of culpable ignorance can be given in which (roughly) a person is culpably ignorant regarding x at time t only if most people in relevantly similar circumstances would have known x at t . Rather than focusing on prior facts about the agent in question, as a tracing-based approach would do, such an account locates culpability in extra-agential facts. Though it is not my intention

²⁴ A referee suggested that IEC is "to a certain extent, ... *ad hoc*, as it is explicitly aimed at solving certain particularly problematic cases." First, note that IEC is aimed at resolving entire classes of problematic cases, not just particular cases. Second, if refining a condition in the light of whole classes of raised counterexamples makes that condition *ad hoc*, then very few conditions in analytic philosophy would avoid the charge.

²⁵ See footnote 13 above.

²⁶ I am thankful to Manuel Vargas for pushing me to directly address the worry about circularity here and for helping me think about ways that the circularity might be avoided.

here to defend such an account, it is not obviously the case that such an account is a non-starter; and so it is not obviously the case that providing an account of the nature of culpable ignorance presupposed in IEC will be circular.

Second, IEC will be satisfied in a wider range of cases than it otherwise would given that it does not require the agent to have occurrent belief about the moral status of her actions to be morally responsible for those actions. Merely not thinking about the moral status of one's actions ought not undermine one's moral responsibility for those actions. This is true of both non-tracing cases and tracing-based cases. Just as an agent's not having the occurrent belief that *this particular act of lying is morally wrong* doesn't excuse the agent for that particular lie, so too the agent's not having the occurrent belief *becoming a chronic liar might lead me to lie at some future point* doesn't by itself excuse the agent for engaging in a future lie. It is at this point, however, that we directly encounter the worry about what an agent *should* believe about the possible future results of her present actions or patterns of action—and this is, I think, at the heart of the issue regarding tracing.

It is not surprising that giving a detailed answer to this question is difficult, if not impossible. But IEC does give us some guidance. According to (ii) of IEC, what it is reasonable to require that the agent believe about what the future outcomes of his present behavior might be is context sensitive. As the discussion of the fourth desideratum in § 6 showed, the more serious a possible moral consequence is, the less certainty or likelihood required to underwrite the agent's moral responsibility. The kinds of cases that Vargas presents all involve aspects of one's moral character (i.e., character traits, dispositions that are less robust than full-blown character traits, interaction among one's character traits, and the consequences of interaction among one's traits).²⁷ But insofar as the moral characters we develop, and the possible outcomes resulting from the development of those moral characters, are very serious moral issues, moral agents have an obligation to consider future results that, even if not likely, could very well result from our moral characters. To return to a previous example, even if one pattern of action carries with it only a .05 probability that I will become the kind of individual who is insensitive to publication deadlines and another pattern of action carries with it the same probability that I will become the kind of individual who abuses children, the moral significance of the latter should lead me to consider the possibility involved more carefully.

Finally, many cases in which conditions (i) and (ii) of IEC are *not* satisfied will not undermine the agent's moral responsibility in virtue of condition (iii). According to (iii), concomitant ignorance will not function as a moral excuse. Concomitant ignorance involves a certain sort of counterfactual stability with respect to the agent's behavior given that the ignorance plays no role with respect to the agent's performing the action in question. For example, even if an agent

²⁷ I also find it interesting that those cases which involve the less foreseeable interaction among traits (i.e., Paulina the Paralyzed and Ruben the Unfortunate) are also the two cases with the highest level of affect involved (involving an alligator's eating of a child and an inebriated individual causing harm to a child, respectively). For an interesting discussion of the relationship between affect and judgments of moral responsibility, see Nichols and Knobe 2007.

was not able to know that her present course of action involving telling repeated lies would make her into a chronic liar, if such knowledge would have been irrelevant to her freely chosen course of action, there is no reason to allow this ignorance to excuse those lies that later result from her character. And if this is correct, then counterfactual stability of this sort will also preserve the agent's moral responsibility in a range of cases.

IEC also has the virtue of explaining what it is that is effective about Fischer and Tognazzini's Denial Strategy, introduced above in section 4. This is because Fischer and Tognazzini's response to Vargas' cases is bolstered by considerations of IEC.²⁸ I begin with their treatment of Jeff the Jerk. Fischer and Tognazzini claim that in order for a case to be troublesome for tracing, the following three conditions must hold:

- (i) The agent is morally responsible for the relevant outcome
- (ii) The agent does not satisfy the control condition on moral responsibility (at the time of his behavior) with respect to the relevant outcome, and
- (iii) There is no suitable prior time at which the relevant outcome was reasonably foreseeable for the agent (Fischer and Tognazzini 2009, 533).

Fischer and Tognazzini agree that Jeff is morally responsible for the way in which he fired his employees, and thus satisfies (1). They then suggest that Jeff may, in fact, satisfy the control condition on moral responsibility at the time of his behavior. (Fischer and Tognazzini 2009, 533). However, for the sake of argument, they concede that Jeff does satisfy (ii) in order to show that (iii) is not satisfied by Jeff. Fischer and Tognazzini differentiation between three different ways of specifying Jeff's jerky behavior:

- (Outcome1) Jeff fires *those* employees who work for *that* company on *that* precise day in *that* precise manner.
- (Outcome2) Jeff fires some of his employees at some company or other at some point in the future in a despicable manner as a result of his jerky character.
- (Outcome3) Jeff treats some people poorly at some point in the future as a result of his jerky character (Fischer and Tognazzini 2009, 538).

According to Fischer and Tognazzini,

²⁸ I focus only on their response to two of Vargas' four cases. I do not consider their response to the case of Britney the bride because Fischer and Tognazzini think that "[i]n order for this case to have any bite, Vargas first needs to secure both the judgment that she is morally responsible for accepting the proposal *and* the judgment that she didn't satisfy any plausible control condition on moral responsibility at the time of acceptance. But we don't see how Britney's case can be described so as to elicit both of these judgments" (Fischer and Tognazzini 2008, 541). Furthermore, I do not consider here their response to the case of Ruben the Unfortunate, though my discussions of the cases of Jeff the Jerk and Paulina the Paralyzed should make it fairly clear how Ruben also satisfies IEC.

We agree, of course, that it's unreasonable to expect a 15-year-old to be able to foresee what sorts of specific actions he will perform some ten years or more down the road while engaged in the daily routine of a career that he has no idea at this point he is even interested in pursuing. This would certainly set the epistemic bar too high. But we also see no reason to suppose that this is the only way to ground the tracing in question.... While it's extremely implausible to expect Jeff's 15-year-old self to have been able to foresee Outcome₁, it doesn't seem at all implausible to suppose that Jeff's 15-year-old self could have foreseen Outcome₃ (Fischer and Tognazzini 2009, 537).

This much seems right—just as one need not be aware of all the features of an action for which one is morally responsible, one also needn't be in a position to know ahead of time all the specific features of a subsequent action that an action (or course of action) will lead him to perform in the future in order to have the responsible for the later benighted action trace back to the earlier benighting action (or course of action). Fischer and Tognazzini continue:

We suggest that all tracing requires in this case is that Jeff could have reasonably foreseen Outcome₃ at the time he decides to acquire a jerky character. Why do we hold Jeff responsible for unreflectively firing his employees in such a despicable manner? We hold him responsible partly because he freely decided to become a jerk at some point in the past, and it is reasonable to expect Jeff's younger self to have known that becoming a jerk would in all probability lead him to perform jerky actions.... Jeff is indeed morally responsible for his jerky firing of the employees; after all, he could at that "suitable" past time reasonably foresee that his behavior could lead to jerkiness, and he need not be required to envisage the particular details. Thus Jeff can be deemed responsible, and there is no trouble for tracing here (Fischer and Tognazzini 2009, 538).²⁹

Here, Fischer and Tognazzini are claiming that it was not unreasonable for Jeff to be aware that fostering a jerky disposition might lead him to perform jerky actions. And if we measure Jeff against the epistemic condition on moral responsibility in IEC, we can see exactly why this is so. Had 15-year-old Jeff been aware that his present course of action would likely lead him to later engage in jerky actions as a result of his jerky character, that awareness either would have made a difference to his engaging in that pattern of action or it wouldn't. If it were the case that he would have engaged in the same pattern of behavior even if he had been aware of its later effects, then his ignorance is counterfactually irrelevant to his behavior. That is, Jeff would have been concomitantly ignorant. But in this case, Jeff satisfied condition (iii) of IEC, and his ignorance does not function as a

²⁹ As Fischer and Tognazzini note, Vargas concedes that the case of Jeff the Jerk might fail for just this reason (Vargas 2005a, 277). In footnote 12, Vargas attempts to rescue this case by refining its presentation. I agree with Fischer and Tognazzini that this refinement doesn't work.

moral excuse. On the other hand, suppose that 15-year-old Jeff would not have undertaken to become a jerk had he been aware that his doing so would likely lead him to later engage in jerky actions. Surely it is not unreasonable to hold that even 15-year-olds know that treating people in unnecessarily rude and degrading ways is morally wrong. In that case, Jeff satisfies condition (ii) of IEC:

At the age of fifteen, Jeff

- (a) had a reasonable though again not necessarily occurrent belief that treating people in unnecessarily rude and degrading ways is morally wrong;
- (b) had a reasonable though not necessarily occurrent belief that his becoming a jerk might lead to his treating people in unnecessarily rude and degrading ways in the future;
- (c) freely became a jerk; and
- (d) his freely becoming a jerk starting at the age of 15 is what led to his treating others in unnecessarily rude and degrading ways in the future.

In other words, I claim that Fischer and Tognazzini's response to Vargas' Jeff the Jerk case can be grounded in Jeff's satisfying IEC. Fischer and Tognazzini's denial of a counterexample to tracing succeeds because Jeff is an instance of the more general epistemic condition on moral responsibility spelled out earlier.

In their discussion of the case of Paulina the Paralyzed, Fischer and Tognazzini initially claim that it fails to be a case tracing responsibility back to earlier free actions:

If the case is described with an emphasis on the paralyzing role played by her fear and surprise, we will be inclined to judge her not responsible, and if the case is described with an emphasis on her responsibility, we will be inclined to judge that her fear was not in fact overwhelming and perhaps her omission did issue from a moderately reasons-responsive mechanism after all. So, either we will judge Paulina responsible and in control or else not in control and thus not responsible (on any plausible interpretation of control). Either way, the case does not prove troublesome for tracing, since tracing need not be invoked at all (Fischer and Tognazzini 2009, 564).

However, in order to evaluate whether or not the case of Paulina the Paralyzed raises trouble for tracing, they are willing to grant for the sake of the argument both that Paulina is morally responsible for not doing more to save Paul and that she is not free to save him at the time that the alligator moves in his direction. In such a case, tracing must be invoked:

Is there no plausible prior time to trace back to? Vargas seems to suppose that the most natural way to accomplish tracing in this case is to trace Paulina's responsibility back to a time when she could have undergone therapy to rid herself of her fear of toothy reptiles. He rightly

points out, however, that the times when she could have undergone therapy are also times at which she could not have been expected to foresee being in a situation where her fear would play a role in the death of her son. So, this sort of tracing would violate the epistemic condition on moral responsibility.

But we think that, insofar as we are willing to judge Paulina morally responsible at all, there is a much more plausible place to which to trace back. There was no doubt some point prior to her trip to Florida when Paulina could have educated herself about the potential dangers of Florida wildlife. Vargas points out that it was her first trip to Florida and “since she just never bothered to think about it, she did not have any expectations about the native fauna.” But surely this was negligent of Paulina (Fischer and Tognazzini 2009, 543).

As with their response to Jeff the Jerk, Fischer and Tognazzini deny that the case of Paulina the Paralyzed shows that tracing-based accounts of moral responsibility are problematic. And, also like their response to Jeff the Jerk, their response to the case of Paulina the Paralyzed can be grounded in her satisfying IEP, more specifically condition (ii) of IEC:

Prior to their trip to Florida, Paulina

- (a) had a reasonable though not necessarily occurrent belief that her allowing Paul to be attacked by an alligator if she could have prevented it is morally wrong;
- (b) was culpably ignorant of the fact that taking Paul to Tallahassee without first learning about the dangers posed by the native fauna and taking suitable precautions might lead to Paul to be attacked by an alligator in such a manner that she could have prevented by taking suitable precautions;
- (c) freely took her son to Tallahassee without first learning about the dangers posed by the native fauna and taking suitable precautions; and
- (d) her taking Paul to Tallahassee without first learning about the dangers posed by the native fauna and taking suitable precautions leads to her allowing Paul to be attacked by an alligator in such a manner that she could have prevented by taking suitable precautions.

As with the previous case, Fischer and Tognazzini’s denial that the case of Paulina the Paralyzed is problematic for tracing-based approaches to moral responsibility can be grounded in careful attention to the proper epistemic condition on moral responsibility.

Earlier I suggested that one could respond to the purported troubles for tracing via either of two strategies: the Denial Strategy and the Bring it On Strategy. Now we are in a position to see that one can systematically ground the Denial Strategy in the Bring it On Strategy. The revised epistemic condition IEC justifies

Fischer and Tognazzini's denial that Vargas has presented cases which are troublesome for tracing. While I am not so optimistic to think that the above discussion will address every possible tracing case, I think that it does go a significant distance to systematically lessening the threat to tracing-based approaches generated by attention to the epistemic condition.³⁰

Bibliography:

- Dennett, Daniel.** 1984. *Elbow Room: The Varieties of Free Will Worth Wanting*. Cambridge, Mass: Bradford Books.
- Dressler, Joshua.** 2001. *Understanding Criminal Law*, 3rd ed. New York: LexisNexis.
- Feinberg, Joel.** 1987. *Harm to Others*. New York: Oxford University Press.
- Fischer, John Martin, and Mark Ravizza.** 1998. *Responsibility and Control: A Theory of Moral Responsibility*. Cambridge: Cambridge University Press.
- Fischer, John Martin, and Neal A. Tognazzini.** 2009. "The Truth About Tracing." *Nous* 43. 3: 531-56.
- Ginet, Carl.** 2000. "The Epistemic Requirements for Moral Responsibility." *Philosophical Perspectives* 14: 267-77.
- Guerrero, Alexander.** 2007. "Don't Know, Don't Kill: Moral Ignorance, Culpability, and Caution." *Philosophical Studies* 136: 59-97.
- Haji, Ishtiyaque.** 1998. *Moral Appraisability: Puzzles, Proposals, and Perplexities*. New York: Oxford University Press.
- _____. 2008. *Incompatibilism's Allure: Principal Arguments for Incompatibilism*. Ontario: Broadview Press.
- Kane, Robert.** 1996. *The Significance of Free Will*. Oxford: Oxford University Press.
- McKenna, Michael.** 2008. "Putting the Lie on the Control Condition for Moral Responsibility." *Philosophical Studies* 139: 29-37.
- Nichols, Shawn, and Joshua Knobe.** 2007. "Moral Responsibility and Determinism: The Cognitive Science of Folk Intuitions." *Nous* 41.4: 663-85.
- Ross, W.D.** *Foundations of Ethics*. 1939. Oxford: Clarendon Press.
- Smith, Holly.** "Culpable Ignorance." *Philosophical Review* 92 (1983): 543-72.
- van Inwagen, Peter.** 1989. "When Is the Will Free?" *Philosophical Perspectives* 3: 42-53.
- Vargas, Manuel.** 2004. "Responsibility and the Aims of Theory: Strawson and Revisionism." *Pacific Philosophical Quarterly* 85.2: 218-41.
- _____. 2005a. "The Trouble with Tracing." *Midwest Studies in Philosophy* 29: 269-91.
- _____. 2005b. "The Revisionist's Guide to Responsibility." *Philosophical Studies* 125.3: 399-429.
- _____. 2007. "Revisionism." In *Four Views on Free Will*, eds. John Martin Fischer, Robert Kane, Derk Pereboom and Manuel Vargas, 126-65. Malden, MA: Blackwell.
- Zimmerman, Michael J.** 1997. "Moral Responsibility and Ignorance." *Ethics* 107.3: 410-26.

³⁰ I am thankful to Joe Campbell, Tyler Hower, Neal Tognazzini, and two anonymous referees for comments on earlier drafts of this paper. I would like to thank, in particular, Manuel Vargas for his original paper and for pressing these issues in numerous conversations and Alicia Finch for many helpful discussions regarding these issues over the last few years. Previous versions of this paper were presented at Boise State University and the 2007 CRAP (Conference on Responsibility, Agency, and Persons) held at the University of San Francisco. On these occasions, I benefited greatly from comments and questions from John Martin Fischer, Bob Kane, Al Mele, Seth Shabo, Dan Speak, Tamler Summers, Manuel Vargas, Andrew Cortens, and Stephen Crowley.