

Penultimate Draft. Published Version can be found online at
Philosophical Studies:
<https://link.springer.com/article/10.1007/s11098-016-0772-6>

Reasonable Foreseeability and Blameless Ignorance

Daniel Miller

Abstract: This paper draws attention to a fundamental problem for a version of the tracing strategy defended by Gideon Rosen 2004 and Fischer and Tognazzini 2009. I argue that versions of the tracing strategy that require reasonable foreseeability (rather than actual foresight) are in tension with the view that blameless ignorance excuses. A stronger version of the tracing strategy (i.e., one that requires actual foresight) is consistent with the view that blameless ignorance excuses and is therefore preferable for those tracing theorists who wish to continue maintaining that it does.

1. The Tracing Strategy and a Useful Distinction

Ignorance is commonly taken to excuse an agent from blameworthiness for morally wrong behavior or bad consequences that result from it. However, there are cases in which an agent lacks awareness with respect to the moral status of an action, omission, or some consequence thereof and is still blameworthy for it.¹ In such cases we may attempt to explain an agent's blameworthiness for this item by tracing it to his blameworthiness for some prior item that it resulted from. The tracing strategy helps explain how an agent can be blameworthy for some item even though he lacks awareness of its moral status. Consider the following case:

Allergy: Dr. Adams gives Jeremy a certain type of medicine for Jeremy's ailment, unaware that Jeremy is fatally allergic to one of the chemicals in the medicine. Beforehand, however, Adams's

¹ Ignorance with respect to the moral status of an action or omission comes in different forms. In this paper I am concerned with what is often called *circumstantial ignorance*: ignorance of the wrong-making features of an action or omission. I discuss this in more detail in section 4.

nurse asked him if he wanted to look over Jeremy's medical records (which include information about Jeremy's allergy). Adams chose not to look at Jeremy's medical records, aware that medical records often include important information about patients' allergies.

Dr. Adams acts from ignorance in giving Jeremy medicine that he is allergic to. He is blameworthy for this action (and its consequences) despite his ignorance in part because he is blameworthy for something that led to his acting ignorantly, namely, his conscious choice to ignore Jeremy's medical records. However, not just anything that a blameworthy action or omission results in is something for which the agent is thereby also blameworthy. It must also be that, at some time in the past, it was foreseen (or at least foreseeable) that the action or omission would or might result in the subsequent thing. The following case taken from Holly Smith helps illustrate this:

Blood: The daughter of a business executive has been assigned a science project of investigating the distribution of blood-types within the population. To assist her, the executive asks his secretary to survey the office staff and prepare a report on their blood-types by noon. Instead of carrying out this request, the secretary reads a spy novel. At 11:45 a co-worker suffers massive bleeding as a result of an office accident. Paramedics arrive, and could start a transfusion immediately if the worker's blood-type were known. Unfortunately he is unconscious, cannot provide this information, and so dies before the transfusion can be started. Had the secretary followed her boss's orders, she would have known the worker's blood-type and his life would have been saved (1983, p. 505).

Though the secretary may be blameworthy for failing to follow her boss's orders, it would be unfair to blame her for the death of the worker. After all, how could anyone have foreseen that the secretary's failure to follow her boss's orders would result in the worker's death? Contrast this case with *Allergy*. Because Dr. Adams was aware that medical records often include important information about patients' allergies, and that neglecting to look at medical records can result in unwittingly harming patients, Dr. Adams's blameworthiness for giving Jeremy medicine that he is allergic to is traceable to his blameworthiness for choosing to ignore Jeremy's medical records.

Tracing theorists are divided. While some maintain that the cognitive condition on tracing is best formulated in terms of reasonable foreseeability, others contend that it

requires actual foresight.² In what follows I will argue that, for a number of theorists, a commitment to the reasonable foreseeability view is in tension with other views that they maintain about the nature of epistemic obligation and the conditions under which ignorance excuses. The basic structure of the argument is as follows: It's possible that an agent be blameless for failing to foresee what was reasonably foreseeable for him. But, if this is so, then foreseeability views imply that an agent might be blameworthy for some consequence even if he was blamelessly ignorant of it. This implication, however, is at odds with the widely held view that blameless ignorance excuses. Before continuing, I will lay out the tracing strategy and its various formulations in more detail.

The tracing strategy helps highlight a useful distinction. An agent's blameworthiness for something is *direct* when he is blameworthy for it in virtue of the control that he exercises and the awareness that he has at the time. In contrast, an agent's blameworthiness for something is *derivative* when he is blameworthy for it at least partly in virtue of his blameworthiness for some prior thing.³ As I will understand it, the tracing strategy entails the following:

- (TS)** If an agent lacks the control or awareness required to be directly blameworthy for some particular morally wrong or bad item Y, then he is blameworthy for Y if and only if
- (i) Y is (at least partly) the result of some prior thing X for which the agent is blameworthy, and
 - (ii) the agent satisfies the relevant foresight condition with respect to Y.

The second clause in (TS) concerns what I will call the "epistemic condition" on derivative blameworthiness. Conceptual space leaves room for at least four general positions concerning this condition:

- (a) The agent must actually foresee at the time of X that it would or might result in Y.
- (b) The agent must actually foresee at the time of X that it would or might result in a certain type of consequence of which Y is a token.

² Rosen (2008), for example, requires reasonable foreseeability, while Zimmerman 1997 requires actual foresight (pp. 420-421).

³ Zimmerman (1997) and Rosen (2004) give similar accounts of this distinction.

- (c) It must be reasonably foreseeable for the agent at the time of X that it would or might result in Y.
- (d) It must be reasonably foreseeable for the agent at the time of X that it would or might result in a certain type of consequence of which Y is a token.⁴

Corresponding to these four positions on the epistemic condition are four versions of the tracing strategy (TSa, TSb, TSc, and TSd), each differentiated by the particular epistemic condition stated in clause (ii) above.

We can narrow the range of options. In their response to Vargas (2005), Fischer and Tognazzini (2009) convincingly argue that it is an overly demanding epistemic requirement on tracing that particular consequences be reasonably foreseeable for the agent. A good deal of their discussion focuses on Vargas's *Jeff the Jerk* case, in which a teenage Jeff takes on a process of character development (of becoming a jerk) that results in his behaving in a rude and insensitive manner when he lays off employees some decades later. Vargas explains that, although Jeff is intuitively blameworthy for this later act, it was not reasonably foreseeable for Jeff at the time that he began his character development, and so the tracing strategy (construed so as to require reasonable foreseeability) cannot accommodate Jeff's blameworthiness for the later act. But as Fischer and Tognazzini point out, the fact that a *particular* bad consequence of Jeff's becoming a jerk is not reasonably foreseeable for Jeff at the time does not preclude its being reasonably foreseeable that his becoming a jerk will likely result in his behaving in jerky ways down the road. Put simply, they maintain that it is sufficient for tracing that consequences more broadly construed are reasonably foreseeable for an agent. If the correct version of the tracing strategy is one that requires reasonable foreseeability, it is plausibly TSd. This sort of consideration can also be applied to versions of the tracing strategy that require actual foresight; if the correct version of the tracing strategy is one that requires actual foresight, it is plausibly TSb.⁵ I limit the following discussion to these two alternatives, referring to them as foreseeability views and actual foresight views, respectively.

⁴ It might need to be added to these foresight conditions that the agent believes that Y (or some type of consequence of which Y is token) is morally bad or wrong.

⁵ For simplicity, in what follows I will often omit "a certain type of consequence of which Y is a token" when discussing these two epistemic conditions on derivative blameworthiness.

2. Foreseeability Views

Gideon Rosen is committed to a foreseeability view. He writes:

[W]e are not culpable in every case for the bad consequences of our culpable wrongdoing. It would be more accurate to say, as the law sometimes does, that we are on the hook for the “foreseeable” bad consequences of the wrong we do (2008, p. 604).⁶

John Martin Fischer and Neal Tognazzini are committed to a foreseeability view concerning the epistemic condition on responsibility for consequences. The version of the tracing strategy that Vargas (2005) attacks, and that Fischer and Tognazzini (2009) defend, is stated by Vargas as follows:

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

Fischer and Tognazzini accept that, given a broad specification of the outcome, KC correctly states the epistemic condition on responsibility for consequences (pp. 537-538).⁷ I proceed on the assumption that Fischer and Tognazzini maintain that, when conjoined with the fact that an agent has sufficient control (required for blameworthiness) over a prior item that the outcome resulted from, the outcome being reasonably foreseeable is sufficient for both responsibility and blameworthiness for consequences.⁸ I take it, then, that they are committed to a foreseeability version of the tracing strategy as understood above.

⁶ Rosen 2004 is clear that he understands moral culpability as equivalent to blameworthiness (pp. 296-297).

⁷ Although (KC) states a necessary condition and not a sufficient one, Fischer and Tognazzini take it that this is sufficient to satisfy the epistemic condition on responsibility for consequences. This commitment will become clearer in 2.1.

⁸ Although Fischer and Tognazzini (2009) commit themselves to an epistemic condition on responsibility for consequences, they do not explicitly commit themselves to a view concerning the epistemic condition on *blameworthiness* for consequences. In a footnote, they say:

[I]t may well be that you are morally responsible for inadvertently offending someone even though you are not blameworthy for doing so. What sort of epistemic requirement is a condition for blameworthiness is an important and interesting question that we do not take up here.

2.1 Reasonable Foreseeability

These theorists are committed to reasonable foreseeability as the epistemic condition on derivative blameworthiness. But what does reasonable foreseeability amount to? Some notions of reasonable foreseeability are analyzed in terms of what a certain *sort* of agent (e.g., a reasonable agent) *would* foresee.⁹ But whether an agent is responsible or blameworthy for something depends on whether the agent in question has or exercises certain capacities with respect to the item in question (e.g., whether the agent or the mechanism from which he acted was reasons responsive, whether the agent could have done otherwise, etc.).¹⁰ Because of this, the reasonable foreseeability requirement in foreseeability versions of the tracing strategy should be understood as relative to the agent in question. We should ask, not whether something is reasonably foreseeable *simpliciter*, but rather whether something is reasonably foreseeable *for a given agent*.¹¹

Some event Y is a reasonably foreseeable result of some event X for an agent S only if S has the ability to foresee that Y will or might result from X.¹² As Fischer and Tognazzini say of one of Vargas' cases, "tracing can apply simply because it is reasonable to think that the agent *can foresee* outcomes that are broadly specified" (p. 552, my emphasis). Of course, an agent's ability to foresee something often depends upon his ability to *do* certain

Because of this, one might wonder whether Fischer and Tognazzini are committed to a foreseeability view of the tracing strategy as I state it in section 1 (since it concerns blameworthiness).

It might be argued, as Fritz (2014) does, that the position that an agent may be responsible for wrongdoing without being blameworthy for it is unmotivated. If Fritz is correct, then Fischer and Tognazzini cannot plausibly maintain that there are distinct epistemic conditions on responsibility for consequences and blameworthiness for consequences. In any case, given their silence on this issue, the argument I present in this paper may be taken either as a critique of their position (on the assumption that they do in fact hold that these conditions are sufficient for *blameworthiness* for consequences) or else as a reason for them to take a stand against foreseeability versions of the tracing strategy (on the assumption that they are merely agnostic about the issue).

⁹ Husak (2011), for example, maintains that an agent *should* be aware of some risk "if a reasonable person in his situation *would* have been aware of the risk" (p. 207, my emphasis).

¹⁰ See, for example, Fischer and Ravizza (1998), according to which an agent's being morally responsible for an action requires that the mechanism from which the agent acted be moderately reasons responsive.

¹¹ Sher (2009, pp. 100ff) maintains that the standards concerning what a given agent should be aware of are sensitive to facts about that agent's cognitive capacities.

¹² For simplicity, I will often use "Y is reasonably foreseeable for S" as shorthand for "it is reasonably foreseeable for S that Y (or some type of consequence of which Y is a token) will or might result from X."

things the performance of which makes foresight likely for him. For example, Dr. Adams' ability to foresee that the medicine he prescribes to Jeremy might harm him depends upon his ability to check Jeremy's medical history chart for allergies.

It is worthwhile to note that Y's being reasonably foreseeable for S does not entail that it is reasonable to *expect* S to foresee Y, since, depending on S's circumstances, there may not be any reason for S to foresee that Y will or might result from X. For example, although it would be reasonably foreseeable for some other doctor that prescribing a certain medication may result in harm to Jeremy, it would not be reasonable to expect this other doctor to foresee this if he has no reason to prescribe the medicine to Jeremy and isn't intending or considering whether to do so, since in that case he will have no reason to do things to make this foresight likely (e.g., ask Jeremy about known allergies). Since Y's being reasonably foreseeable for S often depends upon S's ability to do certain things the performance of which makes foreseeing Y more likely for S, it would be reasonable to expect S to foresee Y only if S had sufficient reason to do such things, and usually this is the case when S is considering whether to or intending to X or if S has reason to X.¹³ It is reasonable to expect S to foresee Y, then, only if it is Y is reasonably foreseeable for S *and* if S has sufficient reason to consider whether Y might result from some action or event X.

Often the notions of Y being reasonably foreseeable for an agent and of it being reasonable to expect an agent to foresee Y are used interchangeably. For example, of *Jeff the Jerk*, Fischer and Tognazzini write, "there must be a prior time at which Jeff's firing his employees is *reasonably foreseeable*" (p. 536, my emphasis). And, of the same case, they write, "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" (p. 538, my emphasis). Using these two notions interchangeably makes sense when speaking of an agent who Xs, since, if Y is foreseeable for that agent, then he will often also have had sufficient reason to consider the possible consequences of X-ing prior to doing so. In section 3.1 I discuss sufficient reasons agents may have to do such things in terms of moral

¹³ S will also have reason to do such things if S is somehow involved in or obligated to help with another agent's deliberation concerning X.

obligations to improve their epistemic position with respect to the consequences of their actions.

3. The Argument

In this and the following section I reveal a tension between foreseeability versions of the tracing strategy and other views that certain theorists committed to this strategy hold. By offering an argument for the conclusion that foreseeability views are false, I will make perspicuous this tension. The argument has three premises. The difficulty for these proponents of TSd (Rosen and Fischer and Tognazzini) is that they accept two of these and have no principled reason to deny the third. Since the argument is valid, these theorists must either give up their version of the tracing strategy or else give up one of the earlier premises that they are committed to.

Before continuing to discuss it in detail, I will briefly summarize the argument. The argument begins from the premise that it is possible for an agent to be blameless for failing to foresee what was foreseeable for him. The second premise is the principle that an agent is blameworthy for acting from ignorance only if he is blameworthy for that ignorance. If blameless ignorance excuses agents for actions, though, then it also excuses agents for action consequences (the third premise). But, given the first premise, foreseeability versions of the tracing strategy contradict this: they imply that an agent can be blameworthy for some consequence even if he was blamelessly ignorant of it. So, foreseeability versions of the tracing strategy are false.

3.1 Introducing the First Premise

The first premise of the argument concerns the nature of epistemic obligation and the possibility of blameless ignorance. It will be helpful to discuss these issues with reference to a case:

Chocolate: Jon tosses his younger brother's chocolate bar into the next-door neighbor's backyard, where a very mean pit-bull guard dog lives. The pit-bull eats the chocolate bar and

dies as a result (the pit-bull is very sensitive to chocolate). Jon does this freely, knowing that it's wrong to throw his little brother's chocolate bar somewhere that he cannot retrieve it.

Jon is blameworthy for tossing his brother's chocolate bar over the fence. Is Jon also blameworthy for the pit-bull's death? It depends upon further details of the case. Jon is aware that the pit-bull lives in the neighbor's backyard. Jon also knows (though we may suppose that he isn't attending to this at the time) that dogs are, in general, sensitive to chocolate. By considering these facts, Jon could have made the relevant inference that throwing the chocolate bar next door would or might result in harm to the pit-bull; that this consequence would or might result from his action was reasonably foreseeable for Jon at the time of his action. Suppose, though, that Jon failed to foresee this. Does this excuse Jon for the pit-bull's death? Certainly not, if Jon's failure of foresight is something for which he is blameworthy. But what if Jon is *blameless* for this failure?

On foreseeability views, whether an agent's lack of foresight is blameless is irrelevant to whether the agent is blameworthy for the consequence. If such blameless ignorance is possible, then this is problematic for foreseeability theorists who maintain that blameless ignorance excuses (I discuss this in section 4). We can begin the argument, then, with this premise:

(1) Possibly, an agent is blameless for failing to foresee what was reasonably foreseeable for him.

(1) is entailed by the conjunction of two claims: (i) if one fulfills all of one's epistemic obligations with respect to some consequence and yet fails to foresee it, one is blameless for being ignorant with respect to that consequence, and (ii) it is possible that one fulfill all of one's epistemic obligations in order to avoid some ignorance without avoiding that ignorance. Because agents have epistemic obligations to foresee certain reasonably foreseeable results of their own actions, a careful assessment of premise (1) will require clarification of the claim that an agent ought to foresee something.

It is commonly held that moral obligations apply directly only to those things that are under our direct voluntary control.¹⁴ We normally think of obligations as applying to actions or omissions; they concern those actions that we ought or ought not to perform. It's difficult, then, to make sense of what an *epistemic* obligation amounts to without further qualifying it. Beliefs generally do not fall under the category of things that are under our direct voluntary control. So, how are we to make sense of obligations to believe, to know, or to be aware of certain things?

On Gideon Rosen's view, epistemic obligations amount to obligations to *do* certain things that will or might result in an improved epistemic position with respect to one thing or another. Rosen calls these "procedural epistemic obligations." According to Rosen, such an obligation "is not itself an obligation to know or believe this or that. It is an obligation to take steps to ensure that when the time comes to act, one will know what one ought to know" (2004, p. 301).¹⁵ Procedural epistemic obligations may include deliberating about alternative courses of action before performing morally significant actions, seeking advice in cases of difficult choices, paying attention to pertinent features of your situation, reflecting on the permissibility of an action, and so on. As Rosen emphasizes, these obligations are impossible to codify; exactly what epistemic obligations one has in a given situation will depend upon the details of that situation.

In defending the tracing strategy against a concern raised by Angela Smith's case of forgetting a close friend's birthday, Fischer and Tognazzini seem to operate with this view of epistemic obligation:

To see why we think some sort of tracing must be going on in the birthday case, consider the fact that Smith failed to choose to do various things which were such that, had she so chosen, she would have had a better chance of remembering her friend's birthday... Part of what it is to be a good friend is to take these steps to minimize the likelihood that you will forget your friends' birthdays. If you don't take these steps, and then forget, you are legitimately morally assessable for your forgetting precisely because you failed to do something to make your forgetting much less likely (and you were in control of this failure) (2009, p. 550).

¹⁴ There are some who disagree with this, such as Robert Adams (1985) and Angela Smith (2005). However, because this paper is concerned with theorists that do not dispute this claim, I will not attempt defend it here.

¹⁵ See also Rosen (2003, pp. 63-64).

In addition to the implicit agreement with Rosen's view of epistemic obligation, the above passage also seems to imply that an agent who did *not* fail to fulfill any procedural epistemic obligations relevant to his foreseeing Y would be excused for failing to foresee Y. After all, if an obligation to foresee Y is just an obligation to do certain things to help ensure an improved epistemic position with respect to Y, it seems possible for an agent to fulfill whatever procedural epistemic obligations he has with respect to Y without subsequently coming to foresee Y. If this were so, the agent would be blameless for failing to see what was reasonably foreseeable for him and premise (1) would be true. In the following subsection I illustrate this possibility.

3.2 The Case for Possible Blameless Ignorance

Recall Dr. Adams. In a variation of the case, he checks Jeremy's chart, which says that Jeremy is allergic to chemical A. Dr. Adams prescribes B to Jeremy, aware that B is composed of chemicals α , β , and γ . As it turns out, ' γ ' is another name for chemical A. Dr. Adams knows this, but his knowledge of this is not occurrent at the time.

It was reasonably foreseeable for Dr. Adams at the time that prescribing B would result in harm to Jeremy. Dr. Adams had the knowledge and capacity required to make the inference that, since Jeremy is allergic to chemical A, B is partly composed of chemical γ , and γ is identical to A, prescribing B will result in harm to Jeremy. We may further assume that Dr. Adams should have foreseen that prescribing B would result in harm to Jeremy, and this epistemic obligation is a matter of his having had certain procedural epistemic obligations. For example, Dr. Adams had a procedural epistemic obligation to check patients' medical history charts before prescribing them medicine, an obligation that he fulfilled on this occasion. And yet, Dr. Adams failed to foresee that prescribing B would result in harm to Jeremy.

Dr. Adams may have come to realize this if he had first checked with another doctor in his office to make sure that B wouldn't harm Jeremy. But if he did have a procedural obligation to do this, then it would presumably be an instance of a more general procedural

obligation to double-check with another doctor any time (or at least most times) before he prescribes medicine. But it's implausible that he has a general procedural epistemic obligation of this kind. So, it's implausible that he had this obligation with respect to prescribing B to Jeremy. It's more plausible that Dr. Adams had a procedural epistemic obligation to double-check with *himself* by asking himself whether B might be harmful to Jeremy. We can suppose that he did, but simply failed to recall that one of B's components is identical to chemical A (but sometimes goes by a different name).¹⁶¹⁷

This case helps illustrate the claim that it is possible that an agent fulfills whatever procedural epistemic obligations he has to with respect to foreseeing Y without subsequently coming to foresee Y. If an obligation to foresee Y is just an obligation to do certain things to help ensure an improved epistemic position with respect to Y, an agent that fulfills these procedural obligations without subsequently coming to foresee Y is excused (and therefore blameless) for failing to foresee what was reasonably foreseeable for him.¹⁸

Let's return to *Chocolate*. What epistemic obligations might Jon have had in his situation? Since it's plausible that we have a general obligation to check around before

¹⁶ One might point out that Dr. Adams does not merely have procedural epistemic obligations to check Jeremy's medical history chart, to ask himself whether B might be harmful to Jeremy, and so on; he has obligations to do these things *carefully*, and he may have failed to fulfill these more specific procedural epistemic obligations. But the supposition that Dr. Adams can, for example, ask himself whether B might be harmful to Jeremy and yet still fail to realize that it will be harmful to him does not require the assumption that Adams did not take a reasonable amount of time asking himself this question or that he did not ask himself carefully. He may have done so and nevertheless failed to recall what he knew. Although it may be *unlikely* that Dr. Adams fulfill each of these procedural epistemic obligations without coming to foresee that prescribing B would result in harm to Jeremy, there is no reason to suppose it impossible. Since (1) concerns possibility and not likelihood, the possibility of such a scenario is sufficient for its truth.

¹⁷ An additional objection is that, if it's true that Dr. Adams had an obligation to foresee that prescribing B would be harmful to Jeremy, then he had procedural epistemic obligations take whatever steps *would* lead to his becoming aware of this fact. This objection assumes both that there are always steps that would *guarantee* foresight of the relevant fact *and* that an agent is required to take those steps, whatever they happen to be. Each of these assumptions would require support. Furthermore, each assumption is problematic. The first assumption fails to acknowledge that epistemic failures are often to due to circumstantial features of a particular situation in conjunction with the limited cognitive capacities of the agent, so that it is not necessarily true that there are always steps the taking of which would lead to foresight. The second assumption fails to acknowledge the limits on what we can reasonably be expected to do. As Rosen stresses, we are not (legally or morally) required to be *maximally* prudent, but only *reasonably* so (2004, p. 301). I am grateful to an anonymous reviewer from *Philosophical Studies* for suggesting that I address this objection.

¹⁸ It's worthwhile to note that we needn't assume that such an agent's failure is somehow due to a negative quality of will, such as a lack of due regard. On the contrary, often times, an agent who fulfills all of the relevant epistemic duties displays due regard by doing so. Again, I am grateful to an anonymous reviewer from *Philosophical Studies* for suggesting that I address this objection.

throwing something even slightly heavy to make sure that it won't hit someone, Jon plausibly had this procedural epistemic obligation with respect to throwing the chocolate bar. Suppose that he did fulfill this obligation, but that the fulfillment of this obligation didn't lead to foreseeing that harm might come to the dog. Of course, this isn't surprising; Jon's throwing the chocolate bar doesn't pose any serious danger to the dog for the reason that it might hit the dog, but rather because the dog might eat it. This thought might lead one to think that Jon didn't merely have an obligation to check to see whether he might hit anyone. He has a broader procedural epistemic obligation; he has an obligation to think about whether, by throwing the chocolate bar, he might *harm* someone. Suppose he does think about this, but all that comes to mind when he asks himself whether he might *harm* someone is whether he might *hit* someone. That is, he fails to realize at the time that there are more ways to harm someone by throwing the chocolate bar than by hitting someone. We can suppose he would have realized this had he asked himself the same question while throwing a more conspicuously dangerous item, such as a firecracker. In this alternative case, though, Jon may well have additional procedural epistemic obligations that he did not have in the actual case.

Some may take it that Dr. Adams or Jon plausibly had further procedural epistemic obligations than the ones I have discussed, and this might be right. But this wouldn't alter the general point that I have been attempting to illustrate: the fulfillment of one's procedural epistemic obligations in order to avoid some future ignorance does not *guarantee* that one will avoid that ignorance. The considerations raised in this section suggest that it is at least possible (even if it is *unlikely*) that an agent fulfill whatever procedural epistemic obligations he has, and yet for some reason or another his fulfillment of these obligations didn't lead (as we might hope it would) to foresight. If this is so, then the agent is blameless for his failure to foresee what was reasonably foreseeable for him, and so premise (1) is true.

3.3 Objections to the First Premise

A possible objection arises here. It might be claimed that, if an agent fulfills all of his procedural epistemic obligations and yet fails to foresee some consequence of his action,

then the consequence is not reasonably foreseeable for the agent. The reasoning is as follows. On one possible construal of reasonable foreseeability, that Y is reasonably foreseeable for S entails that S *ought* to foresee Y. This, in turn, is just a matter of S's having certain procedural epistemic obligations the fulfillment of which make foreseeing Y more likely for S. If an agent has met all of these, it might be explained, then there is no further obligation to fulfill. Given that there is no further obligation that the agent has not met, then if the agent fails to foresee Y, there just is no sense in which he ought to have foreseen Y. In sum, if "Y is reasonably foreseeable for S" entails that "S ought to foresee Y," then Y is not reasonably foreseeable for the agents I discuss above, and so the cases I discuss fail to establish premise (1).

In response to this objection, we should first consider the claim that must be true in order to refute the argument I offer for (1): if S has no unfulfilled procedural epistemic obligations relevant to Y and S fails to foresee Y, then Y is not reasonably foreseeable for S. This claim is counterintuitive. If I am correct, then it is possible for Dr. Adams and Jon to fulfill all the relevant procedural epistemic obligations and yet fail to foresee. But how would this make it so that these consequences are not reasonably foreseeable for these agents? By stipulation, these agents have the cognitive capacities to foresee what they failed to, and there are procedures (that they are able to carry out) the performance of which makes foresight likely for them. These two facts seem sufficient for reasonable foreseeability.

Furthermore, human cognition is fallible. Rosen himself emphasizes that sometimes we simply fail to notice a relevant fact or fail to foresee something, *not* because we lack the capacity to do so, but simply because our cognitive faculties are not infallible; sometimes our failure to realize some feature of our behavior or to foresee some consequence of our action is simply, as Rosen puts it, "cognitive bad luck" (2008, p. 606-607).

Finally, we can revisit some of the considerations raised in section 2.1. It's implausible that "Y is reasonably foreseeable for S" entails "S *ought* to foresee Y," since agents who have the ability to foresee Y may nevertheless lack sufficient reason to foresee Y. And if they have no reason to foresee, then they are not obligated to.

The above objection depends upon the claim that the epistemic condition on tracing requires that S *ought* to foresee Y. Let us, for the time being, suppose that this is the

proposed epistemic condition. Consider an alteration to the first premise that corresponds to this:

(1') Possibly, an agent is blameless for failing to foresee what he ought to have foreseen.

On the view held by Rosen and Fischer and Tognazzini, epistemic obligations are just procedural epistemic obligations; that is, they are just obligations to perform certain actions. If this is right, then an agent might be excused for failing to foresee Y if the agent failed to fulfill some procedural epistemic obligation relevant to Y and the agent has an excuse for that omission. Because a procedural epistemic obligation relevant to Y is an obligation to do something that will or might result in an improved epistemic position with respect to Y, a failure to fulfill such an obligation may be excused in the same way that other failures to fulfill obligations to perform certain actions are excused; i.e., it may be excused due to either lack of control or lack of awareness.¹⁹ If this is so, then an agent may be excused for failing to foresee what he ought to have foreseen, and so (1') is true.

The objector might point out that, since the agent's ability to foresee Y depends upon his ability to do things that make foresight more likely, a lack of control or awareness with respect to one of these things would make it so that he doesn't have the ability to do these things, and so does not have the ability to foresee Y. If this is so, and if 'ought' implies 'can', then it is false that the agent ought to foresee Y in such a case, and so the case fails to establish (1').

This response from the objector succeeds only if it is impossible for an agent to have an excuse for what he ought to have foreseen. If this is true, then the alteration of the epistemic condition on tracing entails that, if an agent fails to foresee what he ought to have foreseen, then his failure to foresee this is blameworthy. I will argue in section 6 that, given certain plausible assumptions, this epistemic condition collapses into an actual foresight version of the epistemic condition on tracing, and so is unhelpful to the foreseeability theorist.

¹⁹ Although I have not here raised the question of what epistemic conditions there are on direct blameworthiness, a blameless failure to satisfy them would constitute an excuse.

3.3 Commitment to the First Premise

The foregoing considerations lend plausibility to (1): it is possible for an agent to be blameless for failing to foresee what was reasonably foreseeable for him. As we have seen, this point may be disputed. However, it has not been my intention to establish (1) beyond any doubt, but rather to argue for its truth on the view that epistemic obligations amount to procedural epistemic obligations. In addition to this, there is evidence that the foreseeability theorists discussed here accept it.

Rosen is committed to (1). Indeed, it is closely related to a premise in one of his own arguments:

The skeptical argument assumes as a premise that *whenever you are culpable for failing to know some pertinent truth, this is because you are culpable for failing to comply with one of these procedural epistemic obligations...* Take any case of action done from ignorance and let it be stipulated that the agent has been utterly scrupulous in policing his own opinion: he has been as careful and as inquisitive and as reflective as a person in his circumstances should be, and yet he has failed to grasp some crucial fact. I claim that if you bear this stipulation clearly in mind, you will be persuaded in every case that the agent's ignorance is not his fault" (2004, p. 302).²⁰

It follows from Rosen's premise that if one is not culpable for failing to comply with one's procedural epistemic obligations with respect to awareness of some consequence Y, then one is not culpable for failing to foresee Y. The conjunction of this claim with the claim that it is possible for an agent to fulfill all of his procedural epistemic obligations without foreseeing what was reasonably foreseeable for him (as I have argued above) entails that (1) is true.

Fischer and Tognazzini also seem to accept (1). When discussing Smith's birthday-forgetting case, they write:

²⁰ Elsewhere Rosen writes, "We also have a plausible sufficient condition for excusable ignorance, according to which ignorance is excusable when it is nonnegligent— that is, when it persists despite the fact that the agent has taken every required precaution against such ignorance" (2008, p. 604).

Surely if it's true that Smith did everything she could do to make forgetting her friend's birthday less likely, then she isn't responsible if she accidentally forgets anyway. If Smith were to explain to her friend how much care she put into trying to remember the birthday, and that despite her best efforts, she simply forgot, we maintain that her friend would no longer be justified in holding her responsible for forgetting... here we assume that the example can be filled in so that it is clear that all reasonable efforts were made by Smith to ensure that her memory would be triggered... (2009, p. 550).²¹

Here Fischer and Tognazzini seem to accept that Smith would be blameless for failing to foresee some consequence were she to fulfill all of the relevant procedural epistemic obligations, and offer no reasons to think that the consequence is not reasonably foreseeable for Smith. There is good reason, then, to think that they are committed to (1).²²

4. Drawing out the Tension: Blameless Ignorance as an Excuse

The second premise in the argument is a widely held principle about the conditions under which ignorance excuses:

(2) An agent is blameworthy for X-ing from ignorance only if he is blameworthy for the ignorance from which he Xs.

Call this the Blameless Ignorance Principle, or (BI). This principle is maintained by the theorists that I have been discussing (Rosen 2004, p. 300; 2003, p. 64, Fischer and Tognazzini 2011, p. 390), as well as a number of other theorists, including Michael

²¹ Although Fischer and Tognazzini speak in terms of moral responsibility here, elsewhere they maintain that moral responsibility is required for blameworthiness, so they seem to be implicitly committed to my claim (2011).

²² Given what Rosen and Fischer and Tognazzini say, they may hold that, if an agent does everything that can reasonably be expected of him with respect to awareness of some relevant fact and nevertheless fails to become aware of that fact, then the fact is simply not reasonably foreseeable. If this is so, then they may not accept (1). While I defend against this sort of objection in 3.3, it is worth noting that what these theorists say does not immediately commit them to the first premise of the argument. I am grateful to an anonymous reviewer from *Philosophical Studies* for pointing this out.

Zimmerman (1997, p. 411), William FitzPatrick (2008, pp. 601-602), Neil Levy (2009, p. 741), and Holly Smith (2011, p. 118).

(BI) can be understood in at least two different ways, corresponding to two different kinds of ignorance that can figure into an agent's performance of a morally wrong action. First, an agent may be ignorant of the wrongness of his action because he is unaware of the feature(s) of his action that make it wrong. Second, an agent may be aware of the features of the action that make it wrong but unaware that they are wrong-making. The first type is commonly called *circumstantial* ignorance and the second kind *moral* ignorance.²³ While the claim that blameless moral ignorance excuses is somewhat controversial, acceptance of the claim that blameless circumstantial ignorance excuses is widespread.²⁴ Here I restrict the second premise to the claim that blameless circumstantial ignorance excuses.

One compelling reason that (BI), when restricted to circumstantial ignorance, has gleaned considerable support is that ignorance of wrong-making features of an action or omission commonly (if not always) precludes the expression of a negative quality of will, which is commonly thought to be necessary for blameworthiness.²⁵ An agent's blameless ignorance that his action will or might have harmful consequences, for example, commonly precludes the expression of a morally objectionable judgment or attitude with respect to that consequence, since the harm was unintended and unknown.²⁶ Consider, now, a closely related principle:

(BIC): A person S is blameworthy for something Y that is a consequence of S's X-ing out of ignorance (with respect to whether Y would or might result from X) only if S is blameworthy for the ignorance from which he X-ed.

²³ Rosen (2003) also makes this distinction, calling the first type factual (or non-moral) ignorance, and the second moral ignorance. FitzPatrick (2008) and Talbert (2013) both use the term "circumstantial ignorance."

²⁴ Rosen (2003) defends the claim that blameless ignorance of any kind excuses. Talbert (2013) accepts that blameless circumstantial ignorance often excuses but argues that moral ignorance often does not. Clarke (2014) argues that (BI) (understood either in terms of circumstantial or moral ignorance) is lacking in support.

²⁵ Talbert (2013), for example, seems to have something much like this in mind in maintaining that non-culpable *circumstantial* ignorance may excuse an agent from wrongdoing but that non-culpable *moral* ignorance usually does not (p. 226).

²⁶ Levy (2009) offers a distinct reason in favor of (BI), namely that it is unfair to blame someone for doing something if he blamelessly lacked the belief that he has good or compelling reason to omit from doing it (likewise, it is unfair to blame someone for failing to do something if he blamelessly lacked the belief that he has good or compelling reason *to* do it). Blameless ignorance of the wrong-making features of an action (i.e., blameless circumstantial ignorance) may make it so that the agent blamelessly lacks such a belief.

Do we have good reason to accept (BIC)? It might be supported by appealing to (BI), and arguing for the following conditional: if (BI) is true then so is (BIC). Indeed, the reason offered above in favor of (BI) lends support to (BIC), as well. As I mention above, blameless ignorance of the consequences of one's action or omission commonly (if not always) precludes the expression of negative quality of will with respect to that consequence, and this fact supports (BI) and (BIC) alike. An additional reason to accept the conditional claim is that it seems that objections against (BIC) work equally well (or poorly) against (BI), so that the two stand or fall together.²⁷

Perhaps the simplest and strongest support for the conditional claim (and also a straightforward explanation for the two reasons just offered) is that (BIC) seems just to *be* a particularized version of the blameless ignorance principle applied to consequences: circumstantial ignorance with respect to X-ing can include ignorance with respect to the consequences of X-ing. Because of the forgoing considerations, the following premise is quite plausible:

(3) If (BI) then (BIC)

The next step of the argument follows from the second and third premises:

(4) (BIC)

The reader may begin to see how the argument proceeds from here. Given our first premise (that possibly, an agent is blameless for failing to foresee what was reasonably foreseeable

²⁷ A second reason to accept the conditional claim can be seen by considering how one might object to these principles. For example, one may reject (BIC) by maintaining that, if it is reasonable to expect an agent to know or to be aware of some likely bad consequence of some (free) action of his and yet he fails to be aware of it, then he is blameworthy for that consequence, blameless ignorance notwithstanding. The purpose of stating this objection is not to defend against it but rather to point out that the same sort of objection works equally well against (BI): if it is reasonable to expect an agent to know or to be aware of some wrong-making feature of some (free) action of his and yet he fails to be aware of it, then he is blameworthy for performing that action, blameless ignorance notwithstanding. Indeed, it may be that any objection to (BIC) can be made to work equally well against (BI). This would support the claim that, if (BIC) is false then so is (BI), which is logically equivalent to the claim we are seeking support for.

for him), foreseeability views have the implication that an agent can be blameworthy for some consequence even if he was blamelessly ignorant of it. Although this claim does not directly contradict the Blameless Ignorance Principle, it does contradict its corollary, (BIC). We can now state the fifth step of the argument:

(5) If reasonable foreseeability is sufficient to fulfill the foresight condition on derivative blameworthiness, then (BIC) is false.²⁸

Finally, it follows from (4) and (5) that

(C) It is not the case that reasonable foreseeability is sufficient to fulfill the foresight condition on derivative blameworthiness.

It follows from (1) - (5), then, that foreseeability versions of the tracing strategy are false. We can condense the argument as follows: If (BIC) is true and if it is possible for an agent to be blameless for failing to foresee what was reasonably foreseeable for him, then foreseeability views are false. The problem is that the foreseeability theorists discussed in this paper accept two of the three assumptions of the argument (1 and 2), and seem to have no principled reason to reject (BIC). Something has to give.

The first option for such theorists is to reject the claim that it is possible for an agent to be blameless for failing to foresee what was reasonably foreseeable for him (premise 1). But the authors discussed here have strong reasons not to reject (1). It seems that (1) is entailed by the view that we can be directly responsible or blameworthy only for those things that are under our direct voluntary control. On this view we can be directly responsible or blameworthy only for actions or omissions. This “control-based” view is a foundational view that these theorists are committed to about the nature of responsibility and blameworthiness.²⁹ Since epistemic states (ignorance, in particular) are not under our direct voluntary control, we can only be derivatively blameworthy for ignorance by being

²⁸ Notice that both (1) and (5) can be altered to correspond to the amendments to the epistemic condition considered in section 3.2, so that modified versions of this argument will fare just as well against views that adopt those epistemic conditions.

²⁹ Rosen (2004, p.302), Fischer and Tognazzini (2009, pp. 549-550).

blameworthy for something that resulted in it. That is to say, it cannot be that an agent is not blameworthy for any action or omission and yet is nevertheless blameworthy for being ignorant. We can continue this line of reasoning by briefly rehashing some of the motivation offered in section 3 for premise (1). Agents have procedural epistemic obligations to do or to refrain from doing certain things such that fulfilling these obligations makes it likely that the agent will become aware of some morally relevant fact (e.g., that a certain action is wrong, or that it will have bad consequences). But the fulfillment of these procedural epistemic obligations does not guarantee that the agent will avoid the ignorance that the procedural epistemic obligations are aimed at avoiding. This means that an agent may fulfill all of his epistemic obligations without avoiding that ignorance. But then it is possible for an agent to be ignorant and yet blameless with respect to his procedural epistemic obligations. If this is so, and if blameworthiness for ignorance is always derivative (something that seems to be entailed by the control-based view of responsibility stated above), then (1) is true. In 3.2 I point out that these theorists may attempt to resist this line of argument but that doing so requires one or more alterations to their view of the epistemic condition that fare no better with respect to avoiding corresponding versions of the first premise.

There are two remaining options for these theorists. They can either reject (BI), or they can adopt a version of the tracing strategy that is not vulnerable to the sort of argument I have offered against foreseeability views.³⁰ It is clear how rejecting (BI) allows these theorists to avoid the argument presented above, since (BI) is its second premise. It may not be as clear, however, how a different version of the tracing strategy helps avoid the same sort of argument. In the next section I will explain how a similar argument cannot be used to show that actual foresight views are false.

5. A Way Out: What TSb Implies

Here I will illustrate how an actual foresight view deals with blameless ignorance (and thus avoids the sort of argument presented above) by showing how it does so with

³⁰ Zimmerman, for example, maintains (BI) and an actual foresight view much like TSb (1997, pp. 420-421).

respect to Jon in *Chocolate*. Recall the epistemic condition that TSb places on derivative blameworthiness:

- (b) The agent must actually foresee at the time of X that it would or might result in a certain type of consequence of which Y is a token.

According to TSb, Jon is blameworthy for the pit-bull's death only if it is the foreseen result of something for which he is blameworthy. If he fails to foresee that tossing the chocolate bar into the pit-bull's yard will or might result in harm to the pit-bull, then TSb implies that Jon is blameless for the pit-bull's death. Notice, though, that the fact that Jon failed to foresee *at the time of his action* that it would or might result in harm to the dog is not, on its own, sufficient for TSb to imply that Jon is blameless for the dog's death. This is because Jon's harming someone as a result of ignorance might be a foreseen result of some *earlier* item for which Jon is blameworthy. If it were, then Jon would be blameworthy for his subsequent ignorance (with respect to tossing the chocolate bar) and for the death of the pit-bull. According to TSb, then, Jon is blameworthy for the death of the pit-bull only if he is blameworthy for his ignorance that his action (i.e., tossing the bar) might have that result. In this way, TSb avoids the implication that foreseeability views have, i.e., that agents can be blameworthy for foreseeable consequences, blameless ignorance notwithstanding.

Some may object that if TSb were true then agents would have an easy way to avoid blameworthiness. Since (b) requires that an agent have actually foreseen the consequences of his actions in order to be blameworthy for those consequences, TSb implies that an agent could simply fail to consider the foreseeable consequences of his blameworthy actions and in doing so avoid blameworthiness for those consequences. Or at least that is the worry.

This objection rests on an oversimplification of TSb. As I pointed out earlier in this section, the fact that an agent fails to foresee *at the time of his action* that it would or might result in certain bad consequences is not, on its own, sufficient for TSb to imply that the agent is blameless for those consequences. That the agent would or might bring about bad consequences as a result of ignorance might be a foreseen upshot of some other thing for which the agent is blameworthy.

The intuitive pull of the objection rests on a misleading picture: If TSb were true, and an agent knew this, he could seemingly avoid blameworthiness simply by choosing not to consider the foreseeable bad consequences of his blameworthy actions. It's not difficult to see why this is false. The agent in question intentionally omits to consider possible negative consequences for the purpose of avoiding blameworthiness. But, he can have this motivation only if he is aware that his ignorance might result in negative consequences of some type, i.e., those very consequences with respect to which he wants to avoid blameworthiness. Since when he chooses not to consider the possible negative consequences he does so with an awareness that doing so may result in his bringing about negative consequences via ignorance, TSb implies that such an agent *would* be blameworthy for those consequences.³¹

TSb implies that an agent is blameless for unforeseen negative consequences of his blameworthy actions only if his failure of foresight is blameless. The upshot in this section is that the argument offered in this paper against foreseeability views cannot be rephrased to work against actual foresight views like TSb. The argument I have offered against foreseeability views rests upon the following implication: when conjoined with premise (1) (i.e., possibly, an agent is blameless for failing to foresee what was reasonably foreseeable for him) foreseeability views imply that an agent can be blameworthy for some consequence even if he was blamelessly ignorant of it, and so (BIC) is false. This implication does not hold for actual foresight views.

6. *An Attempt to Revise the Foreseeability View*

Some have argued that a commitment to the tracing strategy requires admitting that many agents that we would normally take to be blameworthy are in fact not.³² The reason cited to support this claim concerns the epistemic condition that the tracing strategy places on derivative blameworthiness. The alleged problem is that it is rare that outcomes are foreseen (or even reasonably foreseeable) at the time of a prior act or omission for which

³¹ Referencing Aquinas's *Summa Theologiae*, I-II, Q 6, art. 8., FitzPatrick (2008) calls this kind of ignorance (where one intentionally remains ignorant in order to have an excuse) "affected ignorance."

³² Zimmerman (1997), Rosen (2004), Vargas (2005), McKenna (2008), Sher (2009).

the agent is blameworthy. But, since actual foresight is a more stringent requirement than reasonable foreseeability, the set of cases in which agents satisfy the foresight condition on derivative blameworthiness given TSb is smaller than the set of cases in which agents satisfy the foreseeability requirement given TSd. The result, then, is that if a foreseeability view such as TSd has the result that agents are blameworthy significantly less often than we are inclined to think, an actual foresight view like TSb implies that even fewer agents are blameworthy.

Given the argument I have presented, it might be suggested that the foreseeability theorist can revise the foresight condition without needing to adopt something as strong as the actual foresight requirement set out by (b). With this in mind, consider the following revision of (d):

(d') It must be foreseeable by the agent at the time of X that it would or might result in a certain type of consequence of which Y is an instance, *and if the agent fails to foresee this then his failure must be blameworthy.*

The problem with (d) is that, given premise (1), it implies that agents who blamelessly fail to foresee negative consequences that are foreseeable for them are blameworthy for those consequences, and so TSd implies that (BIC) is false. (d') avoids this problem, so that agents who fail to foresee negative consequences that are foreseeable for them are blameworthy for those consequences only if their failure to foresee (i.e., their ignorance) is itself blameworthy.³³ This does not imply, though, that the foreseeability theorist can avoid an actual foresight view with this revision. Indeed, I will argue here that, given certain plausible assumptions, (d') collapses into (b).

The first assumption is that blameworthiness for ignorance is always derivative (an assumption that the targets of my arguments seem to share).³⁴ If this assumption is true,

³³ The argument in this section works just as well if we add that *S ought to foresee a certain type of consequence of which Y is an instance*" to (d'). The result would be the final revision to the epistemic condition that I consider in section 3.3.

³⁴ This paper is primarily focused on accounts (like Rosen's and Fischer and Tognazzini's) that hold that one is directly blameworthy for something only if it is under one's direct voluntary control. Since ignorance is not under an agent's direct voluntary control, these accounts imply that one cannot be directly blameworthy for

then if an agent is blameworthy for ignorance it is in virtue of being blameworthy for some prior item that resulted in it. Suppose, then, that blameworthiness for ignorance is always derivative and that some agent *S* is blameworthy for his ignorance. Call this ignorance *i*₁. According to (d') *S* is blameworthy for *i*₁ only if it is the foreseeable result of some prior item *X* for which *S* is blameworthy. Suppose that, at the time of *X*, *S* failed to foresee that *i*₁ would or might result from *X*. (d') entails that *S* is blameworthy for *that* ignorance, *i*₂. But, if blameworthiness for ignorance is always derivative, we must again trace *S*'s blameworthiness for *i*₂ to some prior act or omission. As long as any prior act or omission we trace blameworthiness for some ignorance (in this case, the lack of foresight, *i*₂) to is itself accompanied by a lack of foresight (*i*₃) that ignorance would or might result from it, (d') entails that *S* is blameworthy for that lack of foresight (*i*₂), but only derivatively so.

The second assumption is that if *S* is derivatively blameworthy for some instance of ignorance, then *S* is blameworthy for it in virtue of being *directly* blameworthy for some earlier item.³⁵ But, given the conclusion of the preceding paragraph, *mere* foreseeability (without foresight) only yields derivative blameworthiness. The only way that *S* can be blameworthy for any instance of ignorance on (d'), then, is in virtue of being directly blameworthy for some earlier act or omission at the time of which *S* *actually foresaw* that his action or omission would or might result in the ignorance that it results in. But, if this is so, then (d') collapses into (b).

Conclusion

There seem to be only two viable options for the theorists discussed in this paper: they must either give up (BI) or else adopt an actual foresight version of the tracing strategy. I have not attempted to adjudicate between these two options. Rather, I leave these authors with a dilemma. Although it is open to them to reject (BI), tracing theorists who are committed to the view that blameless ignorance excuses have reason to reject foreseeability views in favor of an actual foresight view of the tracing strategy.

ignorance. Rosen (2004, pp. 301-303) explicitly uses this first assumption as a premise in his argument for skepticism about moral responsibility.

³⁵ Rosen also uses this second assumption as a premise in his argument (2004, p. 299).

Acknowledgements: I am grateful to Randolph Clarke, John Fischer, Justin Capes, Kyle Fritz, Gabriel De Marco, Yishai Cohen, and Benjamin Matheson for helpful feedback on earlier drafts of this paper. I am also grateful to the participants at the Moscow State University Free Will and Moral Responsibility Summer School 2014 for comments on an earlier draft of this paper. In addition, I would like to thank Andrew Martin for a thoughtful conversation about this paper.

REFERENCES

- Adams, Robert. (1985). Involuntary Sins. *Philosophical Review*, vol. 94(1), pp. 3-31.
- Clarke, Randolph. (2014). *Omissions: Agency, Metaphysics, and Responsibility*. New York: Oxford University Press.
- Levy, Neil. (2009). Culpable Ignorance and Moral Responsibility: A Reply to FitzPatrick. *Ethics*, 119(4), pp. 729-741.
- Fischer, John and Ravizza, Mark. (1998). *Responsibility and Control: A Theory of Moral Responsibility*. Cambridge: Cambridge University Press.
- Fischer, John and Tognazzini, Neal. (2009). The Truth about Tracing. *Noûs*, 43(3), pp. 531-556.
- Fischer, John and Tognazzini, Neal. (2011). The Physiognomy of Responsibility. *Philosophy and Phenomenological Research*, 82(2), pp. 381-417.
- FitzPatrick, William J. (2008). Moral Responsibility and Normative Ignorance: Answering a New Skeptical Challenge. *Ethics*, 118(4) pp. 589-613.
- Fritz, Kyle G. (2014). Responsibility for Wrongdoing without Blameworthiness: How it Makes Sense and How it Doesn't. *The Philosophical Quarterly*, vol. 64(257), pp. 569-589.
- Husak, Douglas. (2011). Negligence, Belief, Blame and Criminal Liability: The Special Case of Forgetting. *Criminal Law and Philosophy*, 5(2), pp. 199-218.
- McKenna, Michael. (2008). Putting the lie on the control condition for moral responsibility. *Philosophical Studies*, 139, pp. 29-37.
- Rosen, Gideon. (2003). Culpability and Ignorance. *Proceedings of the Aristotelian Society*, 103(1), pp. 61-84.
- _____. (2004). Skepticism about Moral Responsibility. *Philosophical Perspectives*, 18(1), pp. 295-313.
- _____. (2008). Kleinbart the Oblivious and Other Tales of Ignorance and Responsibility. *Journal of Philosophy*, 105(10), pp. 591-610.

- Sher, George. (2009). *Who Knew?: Responsibility without Awareness*. New York: Oxford University Press.
- Smith, Angela. (2005). Responsibility for Attitudes: Activity and Passivity in Mental Life. *Ethics*, 115(2), pp. 236-271.
- Smith, Holly. (1983). Culpable Ignorance. *Philosophical Review*, 92(4), pp. 543-571.
- _____. (2011). Non-tracing cases of culpable ignorance. *Criminal Law and Philosophy*, 5(2), pp. 115-146.
- Talbert, Matthew. (2013). Unwitting Wrongdoers and the Role of Moral Disagreement in Blame. In David Shoemaker (ed.), *Oxford Studies in Agency and Responsibility: Volume 1* (pp. 225-245). Oxford: Oxford University Press,.
- Vargas, Manuel. (2005). The Trouble with Tracing. *Midwest Studies in Philosophy*, 29(1), pp. 269-291.
- Zimmerman, Michael. (1997). Moral Responsibility and Ignorance. *Ethics*, vol. 107(3), pp. 410-426.