

A Puzzle Concerning Blame Transfer

Philip Robichaud & Jan Willem Wieland*

Vrije Universiteit Amsterdam

Suppose that you are a doctor and that you prescribed a drug to a patient who died as a result. Suppose further that you could have known about the risks of this drug, and that you are blameworthy for your ignorance. Does the blameworthiness for your ignorance ‘transfer’ to blameworthiness for your ignorant action in such a case? Many are inclined to accept that such transfer can occur and that blameworthiness for ignorant conduct can be derivative or indirect in this way. In this paper, we motivate a new problem for transfer proponents. Existing accounts cannot satisfactorily distinguish between agents who are blameworthy for their ignorance and their resulting ignorant conduct and agents who are blameworthy for their ignorance but *not* for their resulting ignorant conduct. We defend a solution that we call the ‘concern constraint’ that distinguishes clearly between these cases. Our solution is distinct from the much discussed foreseeability constraint (according to which blameworthiness transfers only when the ignorant action is foreseeable), and it is less restrictive than Holly Smith’s quality of will condition (according to which blameworthiness *never* transfers). We argue that any account of derivative or indirect blameworthiness for ignorant action should take the concern constraint on board.

Keywords: ignorance, blameworthiness, transfer, tracing, mismatch, concern

* This paper was fully collaborative. We would like to thank: two referees of the journal, Ryan Cox, Taylor Cyr, Machteld Geuskens, Alex Guerrero, Liz Harman, David Ludwig, Ben Matheson, Matthé Scholten, George Sher, Jojanneke Vanderveen, and audiences in Amsterdam and Boulder. Thanks also to the Netherlands Organization for Scientific Research for their support through two Veni grants. Email: p.robichaud@vu.nl; jjwwieland@gmail.com

1. The puzzle

*Lazy Doctor.*¹ Julie is a doctor. She is aware that in her specific area she is morally required to spend approximately 10 hours per week keeping up on practice-relevant research, and she heard that a colleague's recent article reports new, important findings about the drug Inscientium. Julie has the time and energy to read it during work hours, but instead chooses to have coffee with a colleague. She knows that she should read the study, but opts for the coffee break out of frustration over how far behind she is on her research. In fact, in the past week she did not spend *any* time reading practice-relevant journals. In the study that Julie fails to read, Inscientium—the most effective drug for treating hay fever—is decisively shown to cause fatal heart attacks in people with rare kidney conditions. The next day, Julie prescribes Inscientium to a patient with one of the rare kidney conditions who suffers a fatal heart attack as a result.

Here are three questions about this case. First: is Julie blameworthy for not reading the study and so for the ignorance from which she acts? Second: is Julie blameworthy for the death of her patient? Third: is Julie blameworthy for the death of her patient *because* she is blameworthy for the ignorance from which she acts? If the answer to all three of these questions is positive, then we have a case of what we call 'blame transfer': Julie's blameworthiness for failing to read the study *transfers* to blameworthiness for the death of her patient. As far as we are concerned, this might well be the right analysis of the case. However, blame transfer is a tricky phenomenon. Compare Lazy Doctor to:

Lazy Friend. Julie is a supererogatory doctor who is extremely concerned about her patients' well-being. She is so concerned that, although she is only morally required to spend approximately 10 hours per week keeping up on practice-relevant research, she always spends 15 hours doing so. Jaap, a science journalist and good friend of hers, is sick and wants to remain up to date during his sickness. He tells Julie about a recent article that reports new, important findings about the drug Inscientium. He asks her to read it and to tell him about its main finding. Julie has already spent 15 hours reading journal articles this week, but she promises to read that extra article to help Jaap. She has the time and energy to do so, but instead she chooses to have coffee with a colleague.

¹ Cases like this have been discussed in the literature since Smith (1983).

She knows this is wrong, but she decides to do it anyway and subsequently fails to update Jaap. The next day, Julie prescribes Inscientium to a patient with one of the rare kidney conditions who suffers a heart attack as a result.

In this case, Julie also seems blameworthy for not reading the study and so for her ignorance. Although certain readers may be skeptical that promises made in the context of friendships have any special moral significance, for our purposes we will stipulate that the moral significance of Julie's failure to read the study rests on the fact that Jaap is depending on her and on the fact that she promised to read it. These two facts suffice for Julie to have at least a *prima facie* obligation to read it. Moreover, it gives Jaap a moral claim against her when she does not know the paper's contents. So, if he were to ask her about them, she would reveal to him that she is ignorant, and she would not seem to have any excuse for failing to know the article's contents.

A further question is whether Julie's blameworthiness for not reading the study transfers to blameworthiness for the death of her patient. In this case, it is less clear that it does. It seems problematic to ground Julie's blameworthiness for the death of her patient on the fact that she was blameworthy for breaking her promise to Jaap. She may be a blameworthy friend, but that does not suffice to make her a blameworthy doctor. Julie should have read the study because of her promise, but that does not make her blameworthy for the death of the patient. Hence, the question arises: *why does blameworthiness seem to transfer in Lazy Doctor, but not in Lazy Friend?* This is the problem that we will address in this paper.

2. Clarifications

Before stating the problem in more general terms, let us explain the difference between these cases in a bit more detail. In both cases, Julie has an obligation to read the study. This is fairly clear in *Lazy Friend*, given that Julie promised she would read the study and that she is fully capable of keeping that promise. The existence of an obligation to read the study may be less clear in *Lazy Doctor*. Why think that doctors have obligations to spend time reading articles, and, more specifically, why think that Julie ought to have read this specific article about *Inscientium*?

Rosen has influentially claimed that agents often have *procedural epistemic obligations* (PEOs) to keep themselves informed about certain topics (2004: 301). Doctors have special PEOs to stay current on the safety of drugs they prescribe. Parents

have PEOs to assess the reputation any of childcare facility they utilize. And all of us, qua moral agents, have PEOs to stay informed about the moral permissibility of our conduct.²

Importantly, as with other moral obligations, there are restrictions on how demanding any PEO can be. A doctor cannot read *every* relevant article if she also wants to actually practice medicine. Thus, even a supererogatory doctor who spends 15 hours per week reading articles knows that there is information in further unread articles that is relevant to her practice. Still, it is implausible to maintain that such a doctor should read these further articles given that it is better to have doctors practicing risky medicine than it is to have them read articles all day. They cannot keep informing themselves forever.³ After all, not giving certain treatments might also be risky. Notably, there is nothing special about this case. Parents, CEOs, drivers, consumers, and others also have PEOs that are only reasonably demanding. There are supererogatory parents, CEOs, and so on, for whom the above features will hold. The puzzle could have been formulated in terms of any of these kinds of supererogatory agents.

This demandingness constraint entails that one has a PEO to take further steps to inform oneself about a certain topic *only if* one has not yet taken a reasonable number of such steps. In real life, it is not so easy to determine when one has taken enough such steps, but in our cases we have made this overly-precise: Julie and doctors like her ought to spend 10 hours per week reading relevant journals. Although this is just a stipulation, you can imagine that this requirement is part of a well-justified, code of professional conduct. If you think 10 hours is too much, please adjust the numbers accordingly. What is important is that Lazy Friend, in contrast to Lazy Doctor, has met her PEOs.

Another idealized feature of the cases concerns the content of the article. As we describe it, the content of the article contains *decisive* evidence that Inscientium is dangerous to a specific set of patients. In reality a single study rarely provides this sort of evidence. Rather, the relevant standards of care are typically shaped by the accumulation of evidence that is suggestive. Of course, this does not rule out the possibility of high-powered, relatively decisive studies so we are happy to stipulate that the Inscientium study is one of them. Moreover, the puzzle we develop does not depend

² We take PEOs to be derivative rather than freestanding obligations (in the terminology of Smith 2014), i.e. PEOs enable one to carry out other obligations one has.

³ This is what Gwen Bradford (2017) calls 'The Life Takeover Worry'.

on the accumulative nature of medical knowledge. The form of the puzzle can be realized with respect to different situations and PEOs.

Still, even if one accepts that doctors like Julie have PEOs, why think that in Lazy Doctor Julie has an obligation to read this particular article? Julie's PEO to do 10 hours of research weekly does not specify how she should prioritize from the hundreds of articles that are relevant to her practice. This is a tricky issue, the full treatment of which is beyond our scope. Roughly, we think that the specification of PEOs is governed in part by salience conditions. In Lazy Doctor, Julie's knowledge that her colleague's recent article contains information about a drug she frequently prescribes increases the salience of this article such that in the context of not having met her general PEO to do 10 hours of research, she is now obligated to read this specific article. That this article rather than another one becomes salient for Julie might seem objectionably arbitrary, but there must be workable way for an agent to prioritize what she should read, and her knowledge about the subject of this specific article seems relevant in this case. Of course, the same salience condition can be met in Lazy Friend—she knows that that the article Jaap asks her to read contains findings about a drug she prescribes. But given that Lazy Friend has already met her general PEO, this knowledge plausibly does not trigger an obligation to read the article now.⁴

The plausibility of this claim can be easily missed. It is tempting to think that, in general, agents have a PEO so long as they know that a certain investigative act is likely to provide information that is relevant to discharging one's other duties. In Lazy Friend, it is tempting to think that her knowledge that the article contains new and important findings about Inscientium just suffices for her to have the PEO to read it, regardless of any other efforts she has made. A moment's reflection, however, shows that this general principle is false. Given the unrelenting pace of publications in medical journals and any given practicing physician's knowledge that many publications contain information relevant to their practice, it cannot be the case that she has PEOs to read each of them

⁴ This is not to rule out the possibility that Julie in Lazy Friend could incur an obligation to read an article if, for example, she were informed that it contains some finding that is very alarming and extremely relevant to her practice. The fact that Julie knows that this specific article contains information that could impact her patients' well-being substantially, might trigger an obligation for her to read it even though she has already met the standard 10 hour requirement. This extra information has the effect of making the new article extremely salient, which then triggers the new PEO. Since, in Lazy Friend, Julie is not privy to this kind of information, there is no reason to suppose that a new PEO is triggered.

before treating any patients.⁵ At some point, it must be permissible (if not obligatory) for physicians like Julie in *Lazy Friend* to stop informing themselves. Just when physicians (or anyone else who has PEOs) are permitted to stop informing themselves is a difficult issue in applied ethics. We have just stipulated that physicians have a PEO to read 10 hours of research per week, though the reader might have more or less demanding standards. However demanding these obligations should be, the structure of the puzzle will remain. In *Lazy Friend*, Julie will have read more than the obligatory number of articles and will, thus, be permitted to stop informing herself for that week. There is more to say about this, but a complete understanding to how PEOs function is not necessary for the arguments that follow.

In this debate, the *scope* of Julie's blameworthiness is also crucial. Julie may be blameworthy for (1) failing to read the study, (2) breaking the promise, (3) failing to know that *Inscientium* is dangerous, (4) not being able to inform Jaap, (5) not informing Jaap, (6) prescribing the drug, and/or (7) the death of her patient. In *Lazy Friend*, we assume that (2) is identical to (1) and that (3) and (4) are their consequences. (5), (6), and (7) while not direct consequences of (1), would not normally have occurred if Julie had read the article, and so we will still regard them as outcomes of (1). When we ask whether blameworthiness transfers, we are asking whether blameworthiness for (1) transfers to blameworthiness for (6) and/or (7).

If blameworthiness transfers, then one is derivatively (cf. Rosen 2004) or indirectly (cf. Zimmerman 1997) blameworthy for (6) and/or (7). All philosophers who subscribe to derivative or indirect blameworthiness assume that there is such a thing as the transfer of blameworthiness and will have to grapple with the issues that we address in the following.

Whether one is originally or directly blameworthy for something depends on the exact conditions of blameworthiness. Importantly, we will stay neutral on the exact sufficient conditions according to which one is blameworthy for (1)-(7). There are many theories about blameworthiness for actions, omissions, consequences, and ignorance that specify such sufficient conditions, and it is not our aim to utilize them here. We do maintain that, in both cases, Julie is *not* originally blameworthy for prescribing *Inscientium*. For example, at the time of prescribing the drug, had she believed that the

⁵ A PubMed search for levothyroxine, which is the most widely prescribed drug in the world, shows that roughly 1,000 articles about the drug are published per year. Even if only a fraction of these contained findings relevant to a given physician's specific practice, that would still amount to many tens of articles just for this one drug.

risk was unjustified, she might be considered originally blameworthy. In such a case, the whole issue of blame transfer would not arise. But, as we understand the cases, Julie does not have such beliefs. At the time of prescribing the drugs, she is unaware of any unjustified risk, and hence the question of blame transfer becomes relevant. In *Lazy Friend*, Julie does not believe she runs an unjustified risk because she knows she has met her PEOs for the week, while in *Lazy Doctor* Julie does not believe she runs an unjustified risk because she fails to make the inference from her belief that she has not met her PEOs to the belief that prescribing *Inscientium* is unjustifiably risky. Such failures to carry out even obvious entailments of our beliefs are commonplace, and when they occur, any blameworthiness for subsequent unwitting wrong actions will be derivative.

So, we have two cases on the table, and it has to be explained why blameworthiness does not transfer in *Lazy Friend* while it does in *Lazy Doctor*. Our aim is to offer such an explanation.⁶ This issue matters beyond merely getting the intuitions right. For transfer might not only increase the *scope* of one's blameworthiness (as just explained), but it might also increase the *degree* of one's blameworthiness, i.e. increase how much blame one deserves.⁷ And if *Lazy Doctor* deserves more blame than *Lazy Friend*, we need a story that justifies this difference.

To generalize the problem, we will adopt the terminology from Smith (1983: 547). Omissions to inform oneself, such as failing to read an article, will be called *benighting acts*. The subsequent ignorant conduct, such as prescribing the drug, will be called an *unwitting act*. Benighting acts maintain ignorance, namely the false belief that the unwitting act (such as prescribing the drug) is permissible or the false belief that the unwitting act does not seriously risk a certain bad consequence (such as the death of the patient). In this paper, when we speak of 'outcomes' of benighting acts, we refer to unwitting acts or their consequences. Blameworthiness for the benighting acts will be shortened to 'B1', and blameworthiness for unwitting acts and their consequences to 'B2'. In these terms, the problem is *how and when B1 transfers to B2*. Call this, 'the

⁶ Still, we invite the reader to check whether our upcoming explanation is compatible with his or her preferred conditions of blameworthiness. If certain accounts are incompatible with our explanation (such as Holly Smith's view that we will discuss in section 8), then the latter poses a *prima facie* problem for the former.

⁷ What transfer exactly involves is an important issue that we address in our Wieland & Robichaud (2017).

Transfer Problem'.⁸ Although we focus on the Julie cases and related others, our aim is to cast light on unrecognized problems with extant accounts of derivative blameworthiness and to defend a novel constraint that not only avoids these problems, but governs all cases of blame transfer.

The plan of the paper is as follows. In the next four sections, we will discuss potential solutions to the Transfer Problem, i.e. four potential constraints on transfer: *foreseeability*, *non-deviance*, *basing*, and *concern*. In the end, we will defend our concern constraint and appeal to it to counter Smith's argument that blameworthiness *never* transfers.

3. Foreseeability

One familiar and well-discussed constraint on the transfer of blameworthiness is *foreseeability*. In the following we ask whether it distinguishes Lazy Friend from Lazy Doctor (our answer will be negative). Vargas formulates a version of the foreseeability constraint as follows: "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" (2005: 274). In terms of the terminology just introduced, this constraint can be put as follows:

Foreseeability constraint. B1 transfers to B2 only if the unwitting act was reasonably foreseeable at the time of the benighting act.

Notice that this constraint does not state that an agent must *actually* foresee that she will perform the unwitting act, and, in this respect, the benighting act may be performed without awareness of the downstream risks of failing to inform oneself.⁹ Vargas appeals

⁸ The term 'transfer' might be misleading. Nothing is supposed to *move* from one place to another. Rather, if B1 transfers to B2, then all necessary and sufficient conditions for B2 are fulfilled (i.e. including B1).

⁹ It may be objected that the foreseeability constraint cannot be part of any account of derivative blameworthiness for unwitting acts given that an agent who, at the time of her benighting act, fails to foresee that she might act wrongly in the future will also be ignorant of the wrongness of her benighting act. It may seem to follow that her blameworthiness for the benighting act must also be derivative, which makes the account in some sense regressive. Our response is two-fold. First, since we will argue that the foreseeability constraint is unable to solve the puzzle of blame

to this constraint in order to argue that it is not always possible to trace intuitively plausible cases of blameworthy actions to some prior action because it was not foreseeable at the time of the prior action that the latter action would ensue. One of Vargas's examples is a jerk named Jeff whose present disposition to act rudely and inconsiderately towards his co-workers and to be unreflective about it was the unforeseeable outcome of a plan hatched during his teen years to develop jerk-like traits as a strategy for attracting attention from his peers. Jeff is given the unenviable task of having to inform some of his co-workers that they are being laid off, and he proceeds to do so in a rude and uncaring manner. The backstory makes clear that Jeff simply could not have foreseen that his plan would eventually result in any particular instance of downstream rude behavior. Given this, Vargas claims that any plausible foreseeability constraint will make trouble for proponents of tracing (or transfer, in our terms).¹⁰

Fischer & Tognazzini (2009: 537) have responded by arguing that the question of whether Jeff meets the foreseeability constraint on depends on how his current jerky actions are described. They distinguish between three descriptions of his current action at increasing levels of generality:

- (1) Jeff fires those specific employees on that precise day in that precise manner.
- (2) 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.
- (3) Jeff treats some people poorly at some point in the future as a result of his jerky character.

Fischer & Tognazzini argue that outcomes (1) and (2) are indeed not reasonably foreseeable but that (3) is. They write:

We suggest that all tracing requires in this case is that Jeff could have reasonably foreseen outcome (3) 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

transfer, we are not committed to its viability as part of the most plausible account of derivative blameworthiness. Second, it is controversial whether original blameworthiness for a benighting act requires awareness of its wrongness (cf. FitzPatrick 2009).

¹⁰ 'Transfer' refers to the same relation as the more familiar term 'tracing': if B2 traces to B1, then B1 transfers to B2 (and vice versa).

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 (2009: 538).

Hence, there is some controversy over how the foreseeability condition should be formulated. Does it require foreseeability of some specific future harm, or merely the foreseeability of increased risk that a harm of the relevant type might occur? Because it seems intuitively plausible that Jeff's current actions are blameworthy, foreseeability could be a necessary condition only when the foreseeable action takes a general description. It is not our task to settle this debate.¹¹ Instead, our question is whether some foreseeability constraint may serve as a solution to the Transfer Problem (i.e. that it would entail that blameworthiness transfers in Lazy Doctor but not in Lazy Friend). As announced, our answer here is negative. Note, first, that the outcomes of Julie's failure to read the article admit of different descriptions:

- (1) Julie causes the death of that specific patient by prescribing Inscientium.
- (2) Julie causes the death of an individual at some point in the future by prescribing Inscientium.
- (3) Julie harms an individual at some point in the future by prescribing Inscientium (or more generally a drug that was discovered to be dangerous).

Outcomes (1) and (2) are not reasonably foreseeable for Lazy Doctor or Lazy Friend. In both cases, Julie did not read the study and, thus, cannot foresee that Inscientium is fatal for certain people (let alone for that specific patient).¹² Since (3) does not specify the kind of harm, it seems reasonably foreseeable in both cases. This is clear in Lazy Doctor. She knows that the article contains information about Inscientium, a drug relevant to her practice, but she still decides not to read it. Moreover, she makes this choice in the context of not having met her PEOs. Thus, she should have read this paper but did not. Given this context, it is tempting to think that her blameworthiness for failing

¹¹ For more on tracing, cf. Timpe (2011), Shabo (2015), Talbert (2016: ch. 5). For more on foresight, cf. Zimmerman (1986), Nottelmann (2007: ch. 13).

¹² One might think that both (1) and (2) are reasonably foreseeable, because Julie heard about the study on Inscientium (while Jeff did not hear anything about his future job). If this is right, then they are reasonably foreseeable in both Lazy Doctor and Lazy Friend, and this will not make the difference between them.

to read the study should transfer to blameworthiness for harming her patient precisely because it was foreseeable that being ignorant of the study's contents poses the avoidable risk of causing such a harm.

Notice, however, that it is also true in *Lazy Friend* that Julie knows that the article contains information relevant to her own practice and that her omitting to read it poses a risk of harming some future patient. The difference of course is that in *Lazy Friend*, Julie is a supererogatory doctor and chooses not to read the article in the context of having already read more than the required number of articles. This time, Julie does not run an unjustified risk at the time of prescribing the drugs.¹³ As we discuss above, we make the plausible assumption that agents are permitted to cease their investigative efforts once their reasonably demanding PEOs are met. PEOs like all most moral obligations are short of maximally demanding. It follows that agents who have met their PEOs may permissibly omit further investigative efforts. Thus, *Lazy Friend*, who has actually exceeded her obligations for the week, may permissibly omit reading more articles. Still, even supererogatory doctors can foresee that some permissibly unread article may contain information that would allow her to practice medicine more safely. An important feature of *Lazy Friend* is that the paper Jaap asks Julie to read is relevant for her own practice as well. It follows that (3) is foreseeable in *Lazy Friend* as well, and is no less foreseeable than it is in *Lazy Doctor*.

For each of the different levels of description of the relevant outcome, there is no difference between *Lazy Friend* and *Lazy Doctor* with respect to what is foreseeable. Therefore, foreseeability is not what distinguishes the two cases, and we will have to look for a solution to the Transfer Problem elsewhere.¹⁴

¹³ As noted before, had she believed that she was running an unjustified risk at the time of prescribing the drugs, she might be considered originally blameworthy for her unwitting act.

¹⁴ This also shows that foreseeability, even if necessary, cannot be *sufficient* for transfer of B1 to B2. For in *Lazy Friend*, Julie can foresee that some bad consequence might follow from her omission to read the article, but she does not seem blameworthy for the death of her patient. This is different from the idea that foreseeability cannot be sufficient for transfer given that certain relevant control conditions also have to be met. After all, the latter might well be met in *Lazy Friend*, and yet there is no transfer. In this respect, *Lazy Friend* may constitute a counterexample to the analysis of similar cases offered by Fischer and Tognazzini (2009: 538–9), in which they argue that a 'trace' succeeds when the control and foreseeability conditions are met. Given that their account is meant to capture conditions of moral responsibility rather than blameworthiness, it is not clear how threatening this case is.

4. Non-deviance

Here is a second suggestion: what distinguishes the two cases has to do with the kind of causal chain that links the benighting act and the resulting unwitting one. In *Lazy Doctor* this chain may be of the right kind, while in *Lazy Friend* it may be of a *deviant* sort. Expressed simply, it seems just a matter of accident that Julie's breaking her promise to Jaap led to the death of a patient. To explain this proposal, consider the following case from Smith (2017):

Tom wants to kill his ex-lover Geraldine. Tom buys rat poison with the intention of killing Geraldine by slipping the poison into her food at a restaurant she frequents. Unbeknownst to Tom or anyone else, this brand of rat poison is not toxic to humans but is highly flammable. Tom slips the rat poison into Geraldine's dish before it is to be carried out to her in the dining room, but the dish is left near the restaurant stove and bursts into flame. The fire results in the restaurant's burning down, causing the death of many patrons including Geraldine.

According to Smith, Tom is not blameworthy for how things turned out (namely, Geraldine's death), even though he is blameworthy for other things (such as his decision and attempt to kill her). Still, he seems to satisfy all relevant conditions for blameworthiness. For example, he was free not to slip the poison into her food, he had bad motives, and he could have foreseen, and indeed did foresee, that slipping it into her food would lead to her death. The special feature of the case is that the causal chain between Tom's motivation to kill her and her death is accidental or deviant. Smith argues that if we want to consider Tom blameless, there must be a non-deviance condition on blameworthiness.

Might a similar analysis apply to *Lazy Friend*? If it did, then we would have a constraint that explains why blame transfers in *Lazy Doctor* but not *Lazy Friend*, and our job would be done. In our view, however, a similar analysis does not apply. To see this, we will first translate Smith's suggestion into a constraint on blame transfer (for the case just mentioned involves a deviant causal link between a motivation and an outcome rather than between a benighting and an unwitting act), and then we argue that it does not actually rule out transfer in *Lazy Friend*. The constraint is this:

Non-deviance constraint. B1 transfers to B2 only if the benighting act gave rise, by a non-deviant causal chain, to the unwitting act.

It is difficult to see whether this applies to Lazy Friend without knowing how non-deviance is to be understood. Smith offers no further definition, though on the basis of her case we see at least two plausible readings of this concept: a metaphysical reading and an epistemic one.¹⁵ On the *metaphysical* reading, the causal chain from A to B is deviant if and only if acts or events such as A do not normally or typically lead to B via some causal intermediary C. For example, rat poison does not normally cause someone's death by catching fire and causing fatal burns. On the *epistemic* reading, the causal chain from A to B is deviant if and only if it was foreseen (or foreseeable) *that* A would lead to B, but not *how* this would happen. For example, Tom foresaw Geraldine's death, but not the way in which it happened.

Turning to our cases of unwitting conduct, note that Lazy Doctor meets the non-deviance constraint on both readings of it. The consequence of ignorantly harming one's patients is a normal or typical consequence of failing to keep up with one's practice relevant research. The causal intermediary in this case is becoming ignorant of the dangers of Inscientium. In addition, as we argued in the previous section, it was foreseeable that Julie's failure to keep up with her research will lead to patient harms, and it was also foreseeable how these harms would obtain—via Julie's poorly informed treatment decision.

The question, now, is whether the same can be said in Lazy Friend. As we will show, levels of description will again be relevant here. In the previous section, we argued that the foreseeability of an outcome depends on the outcome's level of description. Similarly, whether there is any deviance in the causal chain between Lazy Friend's benighting and unwitting acts is will depend on how the actions are described.

As a preliminary matter, it is important to notice that the non-deviance constraint is satisfied as long as some non-deviant causal chain exists between the benighting act and the unwitting act *under some description of the relevant actions*. As a necessary condition on blame transfer, the function of the non-deviance constraint is to rule out cases in which there is a merely accidental causal relation between the benighting and

¹⁵ For a statement of the epistemological reading, cf. Sartorio (2007: 751), who takes her inspiration from a case in Feinberg (1970: 195–6). To be sure, we are glossing over significant issues in this debate, and the concept of non-deviance has been utilized in further ways that are not relevant in the present context.

unwitting acts. Cases in which there is a mere accidental causal relation are, by definition, cases in which there is no non-deviant causal connection between an agent's blameworthy failure to inform herself and the subsequent unwitting act under any description of either of these actions. If there were some description of Tom's action such that it would normally lead to Geraldine's death by fire, then Tom cannot maintain that her death was merely an accidental outcome. Similarly, if there is some description of Julie's benighting act (in either case) such that her unwitting act is a normal upshot of her benighting act, then any claim that her unwitting act is nevertheless merely accidental strikes us as unmotivated. As we saw above, in *Lazy Doctor*, such a description plainly exists; failing to keep up with one's practice relevant research normally causes patient harms via the causal intermediary of becoming ignorant about a drug's dangers. Any claim that Julie might make that it was merely accidental that she harmed her patient should ring hollow. Our question, then, is whether or not *Lazy Friend* bears a stronger similarity to *Lazy Doctor* or to Tom.

Let us take the metaphysical reading of the non-deviance constraint first. On first sight, there seems to be deviance in *Lazy Friend*. Promise-breaking does not usually lead to any kind of harm to patients (even though it may harm the friend, who will at least be disappointed that the promise has been broken). Promise-breaking, at the most general level of description, does not seem to have any typical upshot. The upshots and causal intermediaries that stem from a given instance of promise-breaking will depend on the content of the promise. In *Lazy Friend*, Julie breaks a promise to update Jaap on the article's contents, and a normal upshot of this is that Jaap will not be able to finish his article on *Inscientium*. Another normal upshot is that non-doctors will miss the newest research on *Inscientium* because Jaap's article will be incomplete. But, again, this will not harm any patients and may only frustrate the intellectual interests of Jaap's readers.

Yet, we do not think that this establishes deviance in *Lazy Friend*. There is another description of the benighting act that has Julie's ignorant action as a normal upshot. The content of the promise was to read an article that Julie knew to be relevant to her own practice, and doctors who fail to read articles relevant to their practice may end up harming their patients. Julie's failure to keep her promise caused harm to her patient via her holding a false belief about the safety of *Inscientium*. Since this causal intermediary between her failure to read the practice-relevant article, which constituted the promise-breaking, is not at all abnormal, there is non-deviance in this case after all. Given that the benighting and unwitting acts are connected in a typical way under this

description, and given that one such description is enough for non-deviance (as we argue above), the chain in Lazy Friend should be considered non-deviant.¹⁶

Again, imagine that Julie claimed that it was merely accidental that her failure to read an article that she knew was relevant to her practice would lead to her ignorance about Inscientium and eventually to her unwittingly harming a patient. Indeed, this is exactly the claim we imagined Lazy Doctor making above, and it should ring just as hollow. While it is true that breaking a promise does not typically lead to harming one's patients unwittingly, breaking a promise to read something that you know is relevant to practicing medicine safely does have this undesirable effect as a typical outcome.

The epistemic reading of the constraint yields a similar result, and this can be shown much more economically. In both Lazy Doctor and Lazy Friend, Julie could have foreseen not only that failing to read the article would lead to an increased risk of harm to her patients, but also how this would happen (namely, the risk would increase, because she would be ignorant about certain important findings, such as the side-effects of certain drugs).

Thus, despite initial appearances Lazy Friend and Lazy Doctor do not differ significantly in the nature of the causal chain leading from their benighting acts to the bad outcomes of their unwitting act. The non-deviance constraint also fails to distinguish the two cases.

5. Basing

The next potential constraint on blame transfer hones in on the specific moral failures in the two cases. In Lazy Friend, we may assume that the two moral failures in this case are (i) that it is wrong for Julie to fail to read the study because she made a promise to

¹⁶ One might think that a similar gambit will also be available in Smith's case. There, Tom's motive and Geraldine's death might be connected in a normal way only under an extremely general descriptions of the action and outcome, such as 'acting with the intention to kill a person' and 'the person's dying.' Though there seems to be a normal connection here, note that due to the degree of generality of the level of description of the cause 'acting with the intention to kill a person' there is no way to isolate a causal intermediary to test whether it normally leads to the outcome. Once the action and its outcome are somewhat more specified, as they are in Smith's own analysis, then the set of normal causal intermediaries comes into resolution, and it becomes possible to test for deviance.

Jaap, and (ii) that it is wrong for Julie to prescribe Inscientium because this drug causes avoidable harm. These two claims suggest that there is a mismatch between the species of moral reasons that underwrite Julie's moral failures. Call the reasons that underwrite the former wrong 'friend reasons' and those that underwrite the latter 'patient reasons'. The problem with transferring blameworthiness for the wrongful omission to read the article to blameworthiness for prescribing the drug seems to rest on the mismatch between these reasons. Contrast this with an analysis of Lazy Doctor. In this case, there are also two moral failures, namely, (i) that it is wrong for Julie to fail to read the study because Inscientium causes avoidable harm, and reading the study would have enabled her to see that she should not have prescribed the drug, and (ii) that it is wrong for Julie to prescribe Inscientium because this drug causes avoidable harm. The species of reasons that underwrite both of these failures are patient reasons.¹⁷

Before articulating the new constraint, three clarificatory points must be made. First, when we claim that reasons of particular types underwrite certain wrongs, we accept a view about the relationship between reasons and moral wrongness according to which the wrongness of benighting acts and any relevant unwitting acts is a function of moral reasons. On such a view, an action is wrong when moral reasons of sufficient strength recommend not performing it, and there are only insufficiently strong moral reasons that recommend performing it. We assume throughout this paper that what we call patient reasons, which are considerations roughly to do with promoting a patient's well-being, are of sufficient strength to recommend not prescribing Inscientium. For Lazy Doctor, we assume that these same reasons are of sufficient strength to recommend reading the study about Inscientium. For Lazy Friend, we assume that what we call friend reasons, which are considerations of having to do with the value of keeping one's word (for example), are of sufficient strength to recommend Julie's keeping her promise by reading the article about Inscientium.¹⁸

A second related point is that in Lazy Friend Julie also seems to have patient reasons to read the study. After all, Inscientium causes avoidable harm, and were she

¹⁷ These are all *objective* reasons of why certain actions are wrong in the sense that they do not depend either on what Julie believes or what evidence she has. The wrongness of benighting acts is a thorny issue, and in the following we will assume that objective reasons underwrite both the wrongness of failing to meet one's PEOs, in general, and the wrongness of the failure of certain doctors to read specific articles, if those articles are relevant to her practice.

¹⁸ We remain neutral as to which first-order normative theory best captures these assumptions.

to have read the study, she would have realized that she should not prescribe it. Indeed, one might suggest that Julie's patient reasons to read the study may combine with the friend reasons to read it, which would give her obligation a compound normative foundation. If this is right, the differences between Lazy Friend and Lazy Doctor may be only superficial—their obligations to read about Inscientium would have a partially shared foundation. But, the fact that Julie already discharged her PEOs for the week can affect whatever patient reasons she has to read the study in two different ways. First, as we noted in section 2, it seems plausible that for any agent who has discharged her PEOs, she is morally permitted to omit further related investigations. This permission might in turn be grounded on the relative strength of reasons that the agent has to do other things, such as treat patients or enjoy a coffee break. In Lazy Friend, the friend reasons are of sufficient strength to underwrite the obligation to read the article, while the patient reasons are insufficient to ground such an obligation. According to an alternative picture, Julie's patient reasons to read the article may simply lose their strength completely in a context where she has met her PEOs. It follows from this that Julie does not even have insufficient patient reasons to read the study. What we say below does not hang on one of these accounts being true, given that, according to either account, the wrongness of Julie's benighting act rests on friend reasons rather than patient reasons.

Finally we need to clarify the concept of a reasons-match. We will take the following on board. There is a match between the reasons why the benighting act and the unwitting act of an agent S are wrong if and only if

- (i) S's unwitting act A is wrong because p (some reason), and
- (ii) S's benighting act is wrong because p, and informing herself would have enabled her to see that A is impermissible (for if it is true that p, A is impermissible).

With these clarifications made, we can now formulate a new constraint on transfer:

Basing constraint. B1 transfers to B2 only if the reason why the benighting act is wrong matches the reason why the subsequent unwitting act is wrong.¹⁹

¹⁹ Harman (2011: 459) seems to make a similar suggestion when she writes: "the actions and the beliefs are blameworthy for similar reasons." However, Harman does not motivate this in

Notice that this allows transfer in Lazy Doctor, but not in Lazy Friend; in Lazy Doctor there is a reasons-match, while in Lazy Friend there is a mismatch. In both Lazy Doctor and Lazy Friend, it is wrong for Julie to prescribe Inscientium because this drug causes avoidable harm (i.e. a patient reason). But only in Lazy Doctor is it wrong for Julie to fail to read the study for this reason. In Lazy Friend, the benighting act is wrong because she made a promise to Jaap (i.e. a friend reason). Since patient reasons are the reasons why Lazy Doctor's benighting and unwitting acts are wrong, the basing constraint is met in that case. Such a match does not obtain for Lazy Friend. It seems, then, that we have a solution to the Transfer Problem.

The basing constraint also applies in cases where the wrongness of a given benighting act rests on two or more distinct types of wrong-making reasons, each of which is sufficient to make the act wrong. In a variant of Lazy Friend wherein Jaap is a drug sales rep who will relay the contents of Julie's summary to many other doctors, Julie's obligation to keep her promise to Jaap may be underwritten by patient reasons in addition to friend reasons. If Jaap spreads outdated information about the safety of Inscientium, other patients are likely to be put at risk. This suggests that the wrongness of Julie's benighting act in this variant of the case rests on both friend and patient reasons.²⁰

It follows that a match might obtain given that the patient reasons suffice to make Julie's benighting act wrong, and this result does not seem to be threatened by the presence of sufficient friend reasons to read the study. As we have pointed out, the question of whether the patient reasons alone suffice to make Julie's benighting act wrong will depend on other features of the case, such as whether Julie has already met her PEOs. In any case, it is important to note that it is not a condition on reasons-matches that the wrongness of the benighting act rests *only* on the reasons that make the unwitting act wrong. To get a match in the case just mentioned, for example, it is not required that the wrongness of not reading the article rests *only* on the reasons that make prescribing the drug wrong (patient reasons). Blameworthiness can transfer from the benighting to the unwitting act even in cases where the wrongness of the benighting act is overdetermined.

terms of Lazy Friend-like cases, and, indeed, she denies transfer when she adds: "the actions are not blameworthy *because* the beliefs are blameworthy."

²⁰ We owe this example to Liz Harman.

The following case, though, makes more decisive trouble for the basing constraint:

Mistaken Doctor. Julie thinks that she should read a new article on the dangers of the drug Surditas because it is rumored to cause temporary and mild hearing loss. Although she has not read the required amount of practice-relevant research in the past week and has time and energy to read the article, she decides to have coffee with a colleague. Because she fails to read the article she does not learn that the drug actually causes permanent deafness in patients with rare throat conditions rather than the milder symptoms that were rumored. The next day, she prescribes the drug to a patient with this rare condition who subsequently goes deaf.²¹ If the rumor was that Surditas would cause total deafness, she definitely would have read it. We know this, because a week earlier, she read a different article about a different drug precisely because it was rumored to cause total deafness.

Before considering whether Mistaken Doctor is blameworthy for causing her patient to go deaf, it is worth considering carefully the features that distinguish it from previous cases. The first is that, like Lazy Doctor, she has not done the required amount of weekly research, even though she has time to do so. For this reason, it seems plausible that she has an obligation to read the article about Surditas. Moreover, her awareness that she should read the article makes it plausible that she is blameworthy for not doing so. However, like Lazy Friend, Mistaken Doctor is, nevertheless, quite concerned with doing what she can to prevent serious conditions, such as deafness. The problem is that Mistaken Doctor does not act on this concern at the time of the benighting act due to the credible rumor that the Surditas causes only mild harms. Thus, as with Lazy Friend, there is a certain mismatch between what makes her blameworthy. In this case, Julie is concerned enough about her patients to read articles if she suspects they contain information that would forestall significant harms, such as total deafness, but not concerned enough to read articles if she suspects they contain information that would forestall minor harms, such as temporary and mild hearing loss. In virtue of the former concern, it strikes us as problematic to conclude that Julie is blameworthy for causing deafness.

²¹ We do not assume that it is only bad luck that the rumors she heard about the article were false.

Compare Mistaken Doctor to a counterpart doctor who is exactly the same apart from the fact that she is not at all concerned about preventing significant harms, such as permanent deafness. Even in the face of rumors suggesting Surditas might have such harmful effects, the latter doctor still would fail to inform herself. Both doctors are blameworthy for not reading the article, though they are not similarly blameworthy for harming their patients.

Yet, the basing constraint cannot account for this difference. After all, in Mistaken Doctor, there is no mismatch between the type of reason underlying Julie's obligation to read the study and the type of reason that makes her subsequent unwitting act wrong. The relevant reasons are all patient reasons. The same holds if we move from these general reasons (what we have called patient reasons) to more specific reasons (such as those involving either temporary hearing loss or permanent deafness). For (i) prescribing Surditas is wrong because this drug causes permanent deafness, and (ii) not reading the article is wrong because Surditas causes permanent deafness and reading the study would have enabled her to see that she should not prescribe this drug. So even at this level of description, there is a match in both Lazy Doctor and Mistaken Doctor between the reason why their benighting and unwitting acts are wrong. Hence, the basing constraint is satisfied, and yet there still seems to be a mismatch that should be explained.

6. Concern

Is there any better constraint that rules out transfer in both Lazy Friend and Mistaken Doctor? One alternative reason to think that blameworthiness for not reading the study does not transfer to blameworthiness for harming her patient in Lazy Friend is that the former blameworthiness is grounded in a *deficit of concern* for Jaap rather than in a deficit of concern for her patient. One cannot infer from Julie's blameworthiness for her benighting act that she is unconcerned with the kinds of considerations that make her unwitting act wrong. Indeed, in Lazy Friend Julie has a surplus of concern with respect to the latter. This suggests a slightly different constraint that operates with the notion of concern:

Concern constraint. B1 transfers to B2 only if the benighting act expresses a deficit of concern for the same consideration in virtue of which the unwitting act is wrong.

The underlying thought here is that if blameworthiness is to transfer from B1 to B2, then there must be a match between the kinds of reasons for which the agent shows diminished concern in the benighting act and the kind of reasons that underwrite the wrongness of the unwitting act. This seems to work. First, it rules out transfer in Lazy Friend because Julie's omission to read the article expresses a deficit of concern for her friend, and breaking the promise is not the reason why prescribing the drug is wrong. Second, it allows for transfer in Lazy Doctor because Julie's omission to read the article does express a deficit of concern for her patients' well-being, and the threat that Inscientium poses to this is the reason why prescribing it is wrong. Third, it rules out transfer in Mistaken Doctor because Julie's omission to read the article expresses a deficit of concern for mild hearing loss, but this is not the reason why prescribing the drug is wrong (for, the latter is wrong because it causes total deafness).

Like the basing constraint, the concern constraint is motivated by the idea that transfer of blameworthiness requires a certain match between features of the benighting and unwitting acts. The difference between the two constraints lies in the first relatum of the matched features: namely whether it involves moral aspects of the acts themselves or whether it inheres in some agential dispositions to respond to these features (which is one way of understanding moral concern).²²

Before we discuss some refinements to the concern constraint, consider a potential counterexample to this constraint that capitalizes on yet another way in which the wrongness of the benighting act can be mixed. In a different variant of Lazy Friend, Julie's promise to read the study is explicitly tied to considerations to do with the well-being of patients. In this variant, Julie promises to read the study explicitly in order to assuage Jaap's stated concerns about Inscientium, which might be prescribed to him given his condition. As in the original case, Julie has friend reasons to read the study, given that she made a promise. But, Julie's promise in this case is made for reasons having to do with Jaap's own well-being, and these are just the kind of patient reasons that underwrite the wrongness of prescribing Inscientium. Thus, her failure to read the study in this case may indicate a lack of concern for both friend and patient reasons, and

²² One might prefer the basing constraint over the concern constraint given the difficulty of determining a given agent's concern in specific cases (cf. King 2009: 583–7), but we will not pursue this point here. For more on concern, cf. Smith (1991), Arpaly (2003), and Björnsson (2017).

given the latter, her blameworthiness for failing to read the study might well transfer to her unwitting act of prescribing *Inscientium* to her patients.²³

This variant of *Lazy Friend* is similar to the variant above where Jaap was a drug sales rep, and our treatment of it will follow a similar pattern. Instead of distinguishing between two ways in which the benighting act is wrong, here we appeal to the possibility that an agent can express a deficit of concern for two distinct kinds of considerations. If in failing to read the study, Julie expresses a deficit of concern both for considerations to do with keeping one's word and for considerations to do with Jaap's well-being, then we might conclude, which seems plausible, that blameworthiness could transfer to her unwitting act in virtue of the latter.

Note, however, that as with *Mistaken Doctor*, it is important to characterize Julie's deficit of concern at the appropriate level of specificity. In *Mistaken Doctor*, it is not obvious that the expression of a deficit of concern for causing mild hearing loss implies a deficit of concern for causing total deafness. Similarly, in this last variant of *Lazy Friend*, it is not obvious that Julie's expression of a deficit of concern for Jaap's well-being implies that she also expresses deficit of concern for the general well-being of patients. After all, Julie's supererogation is likely the result of admirable degree of the latter kind of concern. If so, her failure to read the study may show a deficit of concern for a different kind of patient reason than that in virtue of which the benighting act is wrong.²⁴ On this variant of *Lazy Friend*, Julie's deficit of concern expressed by her failure to read the study might be limited to considerations to do with to Jaap's well-being, and thus the concern constraint would not license transfer of blameworthiness to her unwittingly harming her patients. The general lesson to draw here is that care is required in properly characterizing both the deficit of concern expressed by an agent's benighting act and the considerations in virtue of which the unwitting act is wrong.

One might object that we have so far offered no principled way of determining the appropriate level of description of the reasons in virtue of which the unwitting act is wrong. Nor have we offered such a criterion for the appropriate level of description of an agent's deficit of concern. In the foregoing, we sometimes appeal to quite general descriptions of the features in virtue of which the relevant unwitting acts are wrong. In

²³ This example is due to Alex Guerrero.

²⁴ It is of course possible that after Julie finishes her 15 hours of research, she becomes fed up with practicing medicine and is no longer concerned with the general well-being of patients. If this happened, then in breaking her promise to Jaap she might indeed express a deficit of concern that is similar to that expressed by Julie in *Lazy Doctor*.

Lazy Doctor and Lazy Friend, we say that the unwitting act is wrong because 'Inscientium harms patients.' Other times we appeal to more specific descriptions. In Mistaken Doctor we say that the unwitting act is wrong because 'Surditas causes deafness' and not because 'Surditas harms patients.' Similarly, our descriptions of Julie's deficit of concern is sometimes general and other times specific. Since we understand deficit of concern as the lack of concern for the wrong-making features themselves or other considerations, deficits of concern seem to inherit the level of description that the relevant consideration takes. So we claim that Lazy Doctor lacks concern for harming her patients and Lazy Friend lacks concern for keeping promises. Keeping things general allows us to get a match in the former and no match in the latter, and thus the concern constraint is satisfied in the former and not the latter. But when it comes to Mistaken Doctor we go specific. Here Julie has a deficit of concern for causing mild hearing loss which, we say, does not match 'Surditas causes deafness,' which is the wrong-making feature of the unwitting act. But what is the motivation for going specific in Mistaken Doctor? Consider the wrong-making features of the unwitting act generally described: Surditas harms patients. Now if we ask whether Julie in Mistaken Doctor has a deficit of concern for harming patients, it seems the answer is yes. Insofar as she has a lack of concern for harming her patients by causing mild hearing loss, she has a lack of concern for harming her patients. Thus, in this case the concern constraint is not satisfied when the wrong-making features of the action and the lack of concern take a specific description, but it is satisfied when the wrong-making features take a general description. Without some way of determining the appropriate level of description, the concern constraint will lead to conflicting results.

In response, note that it is simply inaccurate to say that Mistaken Doctor has a general deficit of concern for harming patients. This would not allow us to distinguish between Mistaken Doctor, who *has* concern for harming patients by causing severe side-effects, but lacks concern for harming patients by causing mild side-effects, and a different doctor who lacked concern for harming patients by causing side-effects of any severity. Only the latter doctor could be accurately characterized as having a general deficit of concern for her patients. The more specific description of Mistaken Doctor's deficit of concern is the right one because it allows us to make this distinction. Note, then, that if the specific description of her deficit of concern more accurately captures Mistaken Doctor's lack of concern, then there is a mismatch with the wrong-making feature of her unwitting act whatever level of description it takes. If it is specific, then the wrongmaking feature is the fact that Surditas harms the patient *by causing a severe side-effects*, and this is a consideration for which Mistaken Doctor actually does have

concern. If it is general, then the wrongmaking feature is the fact that Surditas harms patients, and we have just argued that Julie does not lack a general concern for harming patients.

The general strategy for applying the concern constraint to any case would be first to accurately characterize the agent's deficit of moral concern. An accurate characterization would allow for discriminations between counterparts who have more or less general deficits of concern. For example, does her failure to perform some inquiry indicate a lack of concern for general moral considerations or does she have a more circumscribed, even idiosyncratic, moral insensitivity? With the accurate characterization of the agent's deficit of moral concern in hand, one can then check the wrongmaking considerations of her unwitting act to see if there is a match at any level of description.²⁵

7. Counterexamples?

Hence we suggest that the concern constraint is a plausible constraint on transfer. In the following, we will discuss potential counterexamples to this constraint. In our discussion we hope to illustrate how our account can be employed in the analysis of different kinds of cases of derivative or indirect blameworthiness (and not just doctor cases). Consider the following case:²⁶

²⁵ It is instructive to see how this works in *Lazy Doctor* as well. There, the most straightforward characterization of Julie's deficit of concern seems to take a general description. It seems odd to maintain that she only shows a narrow deficit of concern for harming her patients by giving them *Inscientium*. If, for example, Julie heard that the side-effects of some other drug were the topic of her colleague's article, then she probably still would have gone to grab a coffee. If this is right, then it seems clear that the concern constraint is met, given that her unwitting act is wrong because of the harms the drug causes. However, if Julie's deficit of concern really was restricted to the harm that *Inscientium* causes (and if she would have read the article had it been about some other drug), then an accurate characterization of her lack of concern would take a more specific description. But, even here, there would be a match, given that a plausible characterization of the wrongmaking features of the unwitting act is that *Inscientium* causes harm. In *Lazy Friend*, the mismatch is evident whether Julie's deficit of concern is described generally or specifically. Whether her deficit of concern is for keeping promises in general, keeping promises to Jaap, or keeping promises about reading articles to Jaap, there will be no match with the harm-based considerations that make her unwitting act wrong.

²⁶ This is due to Ben Matheson.

Good Parent. Clive is a supererogatory parent who is very much concerned about his daughter. He usually goes above and beyond the normal call of parental duty. One evening, the teacher phones Clive in order to let him know that his daughter needs to come in early for vaccination. Clive sees who is calling and declines to answer the phone. He misses the news about the vaccinations, and his children then get sick.

In this case, Clive is ignorant about the vaccination plans, and so he fails to get his daughter vaccinated. It is plausible to think that Clive is blameworthy for this unwitting omission, but if he is, this case constitutes a counterexample to our view. The thought is that Clive's failure to answer his phone was blameworthy and that this blameworthiness seems like it should transfer to his unwitting failure to ensure that his children are vaccinated. However, on one natural reading of the case, there is both a mismatch of basing-reasons and a mismatch of concern. There is a mismatch of reasons because the wrongness of the unwitting act (Clive's failure to take his daughter in for the vaccination) is based on considerations having to do with ensuring his child's well-being, whereas the wrongness of the benighting act is based on considerations to do with answering the phone when someone calls. There is a mismatch of concern given that in failing to answer the phone, Clive displays a lack of concern for his daughter's teacher, whereas the wrongness of his unwitting act, as before, is based on that considerations about his child's well-being. Given these alleged mismatches, both the basing and concern constraints are not met, and blameworthiness should not transfer.²⁷ Yet, *Good Parent* seems to be a case where blameworthiness does transfer. Hence, this seems to constitute a counterexample to our position.

Our response is that, on closer inspection, the case admits of *three* distinct readings, none of which pose any problem to our view. According to the first reading, there actually is a match between the reasons that make Clive's benighting and unwitting acts wrong. According to the second reading of the case, there is no match, and the explanation for this challenges the judgment that Clive is blameworthy for his unwitting failure. On the third reading of the case, blameworthiness does not transfer because

²⁷ Even though the basing constraint has Mistaken Doctor-like counterexamples (i.e. is insufficient to rule out blame transfer in all relevant cases), we think it is still a necessary condition on blame transfer, which is not met in the present case.

there is no reason to think that Clive had to answer the phone in the first place. We will discuss these three readings in succession.

In the first reading, according to which it is false that there is a mismatch, we stipulate that the obligation to answer the phone is a *derivative* obligation rather than a freestanding obligation, and, as such, it is based on the risk posed by violating other obligations. This is a plausible stipulation given that there is nothing about the mere act of answering the phone that could ground a moral obligation to do so. Rather, it is the fact that answering the phone often enables us to discharge other moral obligations. On this reading, then, there seems to be a match—Clive's unwitting act is wrong because he should ensure that his children are vaccinated, and the benighting act is wrong for the same reason. Given that answering the phone when teachers call would enable him to meet this obligation, it follows that it would be wrong for him not to do so. The reasons that underwrite this derivative obligation, the violation of which is Clive's benighting act, match the reasons that underwrite the obligation to ensure that his children are vaccinated. In this case we have a match of the relevant kind. By not answering the phone, Clive shows a lack of concern for the reasons in virtue of which his unwitting act is wrong. Since the concern constraint is met such that Clive might indeed be blameworthy for his unwitting conduct, the case does not constitute a counterexample on this reading of it.

On a second reading of the case, Clive had an obligation to answer the phone that was *freestanding*. One might imagine that there are reasons of kindness and responsiveness that are sufficiently strong to ground an obligation to answer the phone. On such a reading, there would indeed be a mismatch, for the unwitting act is wrong because Clive should ensure that his children are vaccinated, whereas the benighting act is wrong because Clive should be kind (and *not*, as with the first reading, because it would have enabled him to see that he should have his children vaccinated). On this reading, our analysis would be similar to Lazy Friend. We do not think Clive is blameworthy for his children getting sick if this is the *only* story about why he should have answered the phone. That Clive's benighting act expresses a deficient of concern with respect to considerations of kindness, does not suffice to establish blameworthiness for failing to ensure that his children are vaccinated, which is an obligation underwritten by altogether different considerations. As with Lazy Friend, we think that once this mismatch is made clear, the case for transfer seems unmotivated.

On a final reading of the case we take a step back and question the assumption that Clive is blameworthy for not answering his phone in the first place. Why think that there is any obligation, freestanding or otherwise, to answer one's phone? There are

several scenarios in which one might think that Clive simply has no obligation to answer the phone. If the details of the case were such that all correspondence about school takes place via email, then there might be no derivative obligation to answer the phone, as there was in the first reading. Imagine further that Clive also happens to disdain the teacher and that his disdain is based on a particular personal history in which the teacher had never called Clive about this children before and, instead, only called Clive to bully him and remind him of the old days. In such a case, there might not even be reasons of kindness that are sufficiently strong to underwrite the freestanding obligation to answer the phone, as there were in the second reading. If the details of the case were such that there was neither a derivative obligation to answer the phone that was based on considerations of ensuring that his children get vaccinated nor a freestanding obligation that was based on considerations of kindness, then it is plausible to think that there is no wrongful benighting act in the first place and, thus, no possibility of blame transfer.

There is another potential counterexample to the concern constraint that takes the form of a drunk-driving case. It is also a case of intuitive blame transfer, yet it seems that there is no match between the deficit of concern expressed by the agent's benighting act and the considerations that make his later action wrong.

Drunk Tractor. Travis is your standard college freshman back from school on break. In his daily life, he is quite conscientious and is not a risk taker. Back home, he meets up with his high school buddies to go out around town. Along the way, someone scores some beers and heavy drinking ensues. After several hours of drinking, Travis and his friends stumble onto a tractor left out in a field overnight. The keys are in the ignition and Travis, in his intoxicated state, cannot resist one last act of mischief. Travis joyrides around the field, which causes massive property damage for the owner of the tractor.²⁸

Intuitively, Travis is blameworthy for the property damage that he causes, and presumably blameworthy for getting so drunk. In addition, it seems that he is blameworthy for the former because of the latter. The problem for our concern constraint is that on the most plausible interpretation of the case, he does not show a lack of concern for causing massive property damage when he performs the benighting act of getting drunk. Although getting very drunk arguably expresses a deficit of concern for those who might be the victims of more predictable types of drunk actions, it is not clear

²⁸ We owe this example to an anonymous reviewer.

that Travis expresses a deficit of concern for the kinds of reasons that make it wrong to steal a tractor and cause massive property damage. From the perspective of someone getting drunk, one simply does not anticipate that the opportunity to do this degree of damage will present itself, and so it does not seem to fall in what Smith calls the “known risk of the benighting act” (1983: 551). If this analysis of the case is right, the concern constraint is challenged. It seems like Travis’s blameworthiness transfers from his getting very intoxicated to his drunk actions and their consequences, even though the concern constraint is not met.

We think that this counterexample, though suggestive, does not succeed. Note first, that it will not do to respond that when Travis got drunk, he expressed a lack of concern for avoidably causing damage to someone’s property, from which it would seem to follow that the concern constraint is met. There would be a match between this generally described deficit of concern and the wrong-making features of his drunk action also generally described. As we recommend above, to check whether the concern constraint is met, first characterize the agent’s purported deficit of concern at the level of description that permits drawing the relevant distinctions between agents. A crucial part of the counterexample is that Travis is typically conscientious, and so he would not have risked getting drunk if he anticipated that he might wreak the kind of havoc he did.

Perhaps one of his friends is not so conscientious, and would drink to excess even if he knew that doing so might cause him to bring about this type of damage. Suppose that everything as it is described above, except that one of Travis’s friends drives the tractor. Unlike Travis, this joy rider regularly engages in risky behavior and is not fazed by the thought that serious harms might result. Travis’s friend clearly has a more severe deficit of concern than Travis does, and thus we should characterize the deficit of concern at a level of description that captures this difference. This difference is occluded if we understand Travis’s actions as expressing a deficit of concern for avoidably causing damage to someone’s property. To draw the relevant contrast with his much more rash friend, we should say that when Travis got drunk he expressed a deficit of concern for causing only the kind of mild damage to someone’s property that are the more typical result of a night out, such as spilling something on someone’s furniture or knocking over something breakable. But, with this more specified level of description of Travis’s deficit of concern in view, it seems much more problematic to maintain that Travis actually is blameworthy for bringing about the massive damage.²⁹

²⁹ Even though he is not morally blameworthy, one might still think that Travis is legally accountable and that he must compensate for the property damage.

Under no description of Travis's concern, then, does this case stand as a counterexample to the concern constraint.

8. Reply to Smith

So our proposal is that the transfer of blameworthiness depends on whether the concern constraint is met. It is an open question whether this constraint is not only necessary but also sufficient for transfer (perhaps jointly so, together with the other constraints discussed earlier, such as foreseeability and non-deviance). Although many philosophers might be inclined to accept this, Holly Smith has presented challenging arguments for the conclusion that blameworthiness never transfers and that we are blameless for *all* ignorant conduct. In the following, we will address her *argument from quality of will*.³⁰ Smith wrote:

The question then must revolve around the issue of ... whether the motives which give rise to his action are reprehensible or not. But the straightforward answer to this question is 'no.' Of course, it is true that at an earlier time, the time of the benighting act, the agent had a reprehensible configuration of desires—a configuration which typically included a willingness to risk eventual wrong—doing of exactly the sort exemplified in the unwitting act. But the fact that he *earlier* had faulty motives does not show that he *now* has faulty motives. (1983: 559)

In terms of our cases, the reasoning is this. At the time of prescribing the drug, Julie lacks reprehensible desires. She may well have a desire to do the right thing, namely to help her patient. If reprehensible desires (or insufficient concern) is required for blameworthiness, as Smith suggests it is, then Julie is not blameworthy for prescribing the drug. At the time of choosing to have coffee with a colleague, though, she did have reprehensible desires, assuming that it is somewhat reprehensible to have a stronger desire to drink coffee than to keep a promise or to keep up with recent research developments. As Smith says: "The doctor who [prescribes the drug] may still be insufficiently concerned about keeping up with medical advances, but *that* is not what

³⁰ In Wieland & Robichaud (2017), we address a second consideration against transfer considered in Smith (1983, 2017), namely the argument from moral luck.

motivates [her to prescribe the drug].” (1983: 559) Smith’s argument can be reconstructed as follows:

- (P1) For all actions A, S is blameworthy for A only if A resulted from insufficient concern on S’s part.
- (P2) For all A, if S was ignorant that A is wrong, then A did not result from insufficient concern on S’s part.
- (C) Hence: for all A, if S was ignorant that A is wrong, S is blameless for A. [from P1, P2]

This yields the extreme view that blameworthiness never transfers. We think this argument is unsound. We do not wish to deny the condition of blameworthiness in (P1) (as announced in section 2, we stay neutral on this). The problematic premise is (P2), which says that A did not result from insufficient concern on S’s part, even if S was insufficiently concerned about the benighting act that led to A. In our view, this may be true in certain cases, but false in others. Sometimes unwitting acts result from sufficient concern, though sometimes they might well result from insufficient concern. In Lazy Friend Julie has sufficient concern for her patients, but Lazy Doctor does not.

To understand Smith’s position, it is important to make a distinction between ‘distal’ and ‘direct’ motives for an action (1983: 563–4). We agree that transfer is problematic if, following Smith, S is blameworthy for the unwitting A only if A expresses a deficit of concern on S’s part *at the time of A* (which constitutes a direct motive). Our concern constraint differs from this in two ways. On the one hand, it relaxes the requirement given that it does not require that the unwitting A expresses a deficit of concern at the time of A. It merely requires that *the benighting act* that led to A expresses a deficit of concern (which constitutes a distal motive). On the other hand, it strengthens the requirement given that not just *any* deficit of concern will render S blameworthy for A. There must be a match. S is blameworthy for the unwitting A only if the benighting act that led to A expresses a deficit of concern *for the same kind of considerations in virtue of which A is wrong*. Lazy Doctor satisfies this condition, while Lazy Friend does not. Hence, in our view, certain ignorant agents do satisfy a relevant quality of will condition for blameworthiness.³¹

³¹ Since Lazy Doctor shows insufficient concern at t_1 but sufficient concern at t_2 , it might be suggested that she is both blameworthy for the unwitting act, relative to t_1 , and blameless for it, relative to t_2 (while Lazy Friend shows sufficient concern for her patients at both times). Although

It seems that in order for our account to stand as a substantive challenge to (P2), it must provide some motivation for abandoning the temporal constraint that is at the heart of Smith's view. One of Smith's concerns seems to be that without this stringent requirement, an agent's blameworthiness for unwitting acts performed in the absence of a lack of sufficient concern will rest on a set of factors that is objectionably expansive or broad (1983: 569–70). The concern constraint, which amounts to a significant strengthening of the conditions of blameworthiness, should go some way towards assuaging worries about overly expanding the bases for blameworthiness.

Smith (1983: 567–8) is aware that (P2) has the implausible implication that agents are never blameworthy for unwitting acts. She diffuses this by appealing to the idea of the scope of blameworthiness. She argues that any blameworthiness for unwitting acts can be accounted for by increasing the scope of what agents are blameworthy for. Her suggestion is that an agent can be blameworthy for an unwitting act only in the sense that she is blameworthy for her benighting act together with certain of its consequences, but that she is *no more* blameworthy for the relevant consequences. In this sense, then, the near-sighted driver might not only be blameworthy for not taking her glasses, but also for the accident caused by her failure to see. While this scoping strategy has promise, and while it may well be the best way to supplement her view, we think that the story is not complete. To adequately capture the difference between Lazy Doctor and Lazy Friend we need to understand *when* increasing the scope is justified. Given that they are both blameworthy for their benighting act, but only Lazy Doctor is blameworthy for the patient's death, one wonders why scope increases only in Lazy Doctor. Our account explains why this is so. Scope increases only when the concern constraint is met.³²

we do see that this accounts for a certain difference between the two cases, we are not attracted to it. According to this picture, both Julies might be blameless for the unwitting act relative to t_2 , but this is mistaken on our preferred analysis. Apart from this, the suggestion raises complicated issues. For one thing, is S both blameworthy (relative to t_1) and blameless (relative to t_2) at t_1 , or is she merely blameworthy (relative to t_1) at t_1 , and then blameless (relative to t_1) at t_2 ?

³² At one point, Smith (1983: 550–1) defines scope by what the agent takes the risk of the benighting act to be. We would add to this that the risk must be believed to be unjustified. After all, in Lazy Doctor Julie may also believe that she runs a risk (namely, a justified one), though we do not think that this suffices for an increase in scope.

9. Conclusion

In this paper, we have raised the problem of when blameworthiness for one's ignorance transfers to blameworthiness for one's ignorant conduct. In response to this problem, we have defended an additional constraint on transfer (i.e. the concern constraint), which is distinct from the much discussed foreseeability constraint, and less restrictive than Smith's quality of will condition.

Why adopt our concern constraint over Smith's condition? First, our concern constraint does not yield the view that we are *never* blameworthy for unwitting acts. Hence, our view is consistent with the commonsense view that we are at least sometimes blameworthy for our ignorant conduct.

Second, our concern constraint is able to capture the intuitively plausible difference between Lazy Doctor and Lazy Friend (and similar pairs of cases), a difference that is obscured by Smith's insistence on the presence of deficit of concern at the time of the unwitting act. We think this is a serious advantage of our constraint. Both doctors may well be sufficiently concerned when they prescribe the drug, though we still would not want to say that overall both doctors are similarly concerned about their patients, and our concern constraint precisely captures this difference.

Third, the concern constraint offers theorists who are inclined to accept a scoping strategy in order to account for derivative blameworthiness in unwitting act cases a device for determining *when* blameworthiness takes scope over both a benighting act and some subsequent unwitting one.

Finally, the concern constraint illuminates an under-explored aspect of the transfer debate, namely that the question of whether blameworthiness transfers from benighting to unwitting acts depends crucially on the levels of description of the relevant components. We defend a procedure for determining the relevant level of description for both the agent's level of concern and the wrong-making features of unwitting acts.

Hence we propose that *any* account of derivative or indirect blameworthiness for unwitting conduct should take our constraint on board. Though it goes beyond the scope of our project here, one can imagine a pair of analogous uncontrolled action cases (such as classic drunk-driving cases, similar to Drunk Tractor discussed in section 7), that have a similar structure to Lazy Doctor and Lazy Friend. If the concern constraint could also be proven to hold in such cases, then our proposal has important implications for derivative blameworthiness *generally*.

References

- Arpaly, N. 2003. *Unprincipled Virtue. An Inquiry Into Moral Agency*. OUP.
- Björnsson, G. 2017. Explaining (Away) the Epistemic Condition on Moral Responsibility. In P. Robichaud & J. W. Wieland eds., *Responsibility: The Epistemic Condition*, ch. 8. OUP.
- Bradford, G. 2017. Hard To Know. In P. Robichaud & J. W. Wieland eds., *Responsibility: The Epistemic Condition*, ch. 10. OUP.
- Feinberg, J. 1970. Sua Culpa. In *Doing and Deserving. Essays in the Theory of Responsibility*, pp. 187–221. Princeton UP.
- Fischer, J. M. & N. Tognazzini 2009. The Truth about Tracing. *Noûs* 43: 531–56.
- FitzPatrick, W. J. 2008. Moral Responsibility and Normative Ignorance: Answering a New Skeptical Challenge. *Ethics* 118: 589–613.
- Harman, E. 2011. Does Moral Ignorance Exculpate? *Ratio* 24: 443–68.
- King, M. 2009. The Problem with Negligence. *Social Theory and Practice* 35: 577–95.
- Nottelmann, N. 2007. *Blameworthy Belief. A Study in Epistemic Deontologism*. Springer.
- Rosen, G. 2004. Skepticism about Moral Responsibility. *Philosophical Perspectives* 18: 295–313.
- Sartorio, C. 2007. Causation and Responsibility. *Philosophy Compass* 2: 749–65.
- Shabo, S. 2015. More Trouble with Tracing. *Erkenntnis* 80: 987–1011.
- Smith, H. M. 1983. Culpable Ignorance. *Philosophical Review* 92: 543–71.
- 1991. Varieties of Moral Worth and Moral Credit. *Ethics* 101: 279–303.
- 2014. The Subjective Moral Duty to Inform Oneself before Acting. *Ethics* 125: 11–38.
- 2017. Tracing Cases of Culpable Ignorance. In R. Peels ed., *Perspectives on Ignorance from Moral and Social Philosophy*, ch. 6. Routledge.
- Talbert, M. 2016. *Moral Responsibility*. Polity.
- Timpe, K. 2011. Tracing and the Epistemic Condition on Responsibility. *Modern Schoolman* 88: 5–28.
- Vargas, M. 2005. The Trouble with Tracing. *Midwest Studies in Philosophy* 29: 269–91.
- Wieland, J. W. & P. Robichaud 2017. Blame Transfer. In P. Robichaud & J. W. Wieland eds., *Responsibility: The Epistemic Condition*, ch. 16. OUP.
- Zimmerman, M. J. 1986. Negligence and Moral Responsibility. *Noûs* 20: 199–218.
- 1997. Moral Responsibility and Ignorance. *Ethics* 107: 410–26.