

Blame Transfer

Jan Willem Wieland and Philip Robichaud

1. Introduction

*Lazy Doctor.*¹ Julie is a doctor. She is aware that in her specific area she is morally required to spend a few hours per month keeping up on practice-relevant research, and she heard that a colleague's recent article reports new, important findings about 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. In fact, in the past month she did not spend *any* time reading practice-relevant research. In the study that Julie fails to read, Inscientium – the most effective drug for treating hay fever – is decisively shown to cause heart attacks in people with rare kidney conditions. The next day, Julie prescribes Inscientium to a patient with those kidney conditions who suffers a fatal heart attack as a result.

Here are two questions about this case. First: is Julie blameworthy for not reading the study? Second: is Julie blameworthy for killing her patient? If the answer to these questions is positive, and if the answer to the former question is relevant to answering the latter question, which seems *prima facie* plausible in this case, then we have a case of what we call “blame transfer”: Julie's blameworthiness for failing to read the study *transfers* to blameworthiness for killing her patient. Hence, “blame transfer” is intended as a shorthand for a certain relation—to be examined in what follows—between two things an agent is blameworthy for.

In principle, Julie may be blameworthy for (1) not reading the study, (2) not knowing that Inscientium is dangerous, (3) prescribing Inscientium, or (4) killing her patient. Let us assume that (2) is a direct consequence of (1), and that (3) and (4), while not direct consequences, are nonetheless consequences of (1) given that they probably would not have occurred if Julie had read the article. The issue of blame transfer, then,

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

is equivalent to the question of whether, and in what sense, blameworthiness for (1) transfers to blameworthiness for (3) or (4).

If blameworthiness transfers, then S is derivatively (cf. Rosen 2004) or indirectly (cf. Zimmerman 1997) blameworthy for (3) and/or (4). All philosophers who subscribe to derivative or indirect blameworthiness are committed to the existence of blame transfer and will have to grapple with the issues that we address in the following. Whether S is originally or directly blameworthy for something depends on the exact conditions of blameworthiness. Importantly, we will stay neutral on the conditions according to which one is blameworthy for (1)-(4). As the chapters in this book testify, there are various different theories on this, and it is not our aim to appeal to some particular account here. Instead, we ask about what it means to say that blameworthiness transfers in cases such as Lazy Doctor.²

But, we do assume that Julie is *not* originally blameworthy for (3) or (4). Had Julie for example believed that prescribing the drugs carried an unjustified risk, then she would plausibly be originally blameworthy for it, and questions about derivative blameworthiness (and blame transfer) would not arise. But, as we understand the case, Julie does not take herself to be running such risks when she prescribes Inscientium, and so the issue of blame transfer becomes relevant.

In order to state the issue in general terms, we will adopt terminology from Smith (1983). Omissions to inform oneself, such as failing to read an article about the dangers of a drug, will be called *benighting acts*. Any subsequent ignorant actions, such as prescribing dangerous drugs while believing that they are safe, will be called *unwitting acts*. Benighting acts either induce or perpetuate ignorance. Because of her benighting act, an agent comes to falsely believe that the unwitting act does not risk a certain bad consequence (such as killing the patient) and that her unwitting act (such as prescribing the drug) is permissible. In this chapter, when we speak of outcomes of benighting acts, we refer to unwitting acts or their consequences.³ In what follows, blameworthiness for the benighting acts will be shortened to “B1”, and blameworthiness for unwitting acts and their consequences to “B2”.

² We assume that if S is blameworthy, then S deserves whatever reactions happen to be constitutive of blame.

³ One might place certain restrictions on when unwitting acts or consequences really count as outcomes of a certain benighting act. One such candidate constraint is sensitivity: X is an outcome of S’s benighting act A only if, holding fixed the interventions of third parties, X would not have occurred had S not performed A (cf. Fischer & Ravizza 1998, ch. 4).

“Transfer” is the counterpart of the more familiar term “tracing”. If B1 transfers to B2, then B2 traces to B1, and vice versa. Both are suitable terms, though they might be misleading. Nothing is supposed to *move* from one place to another. Rather, if B1 transfers to B2, we take it, then all necessary and sufficient conditions for B2 are fulfilled, and two further conditions hold: a necessary condition: *B2 only if B1*; and an explanatory condition: *B2 (partly) because B1*. In Julie’s case, then, blame transfers only if Julie’s blameworthiness for her benighting act is both necessary and explanatory of her blameworthiness for her unwitting act. The necessary condition is widely accepted.⁴ Julie is not blameworthy for prescribing the drug, if she was blameless for her ignorance that she should not do this. But the necessary condition is insufficient for transfer. After all, it is also a necessary condition on Julie’s B2 that she is alive (for example), but this does not involve any kind of transfer.⁵ A further condition is that Julie is blameworthy for her unwitting act *because* she is blameworthy for her ignorance.⁶ What the explanatory condition exactly amounts to is a tricky question to which we return later on.

Consider the following claim: if Julie is blameworthy for her ignorance that she should not prescribe the drug (and if she has no further excuses for doing so), then she is blameworthy for prescribing the drug. We think this is plausible. Generally, we think it is plausible to think that claims of this form are sometimes true. Exactly when blame transfers will depend on further conditions (such as the absence of other excuses one might have for one’s unwitting act) that we need not go into here. What matters is that, on this view, under certain conditions B1 transfers to B2 (call this view “Transfer”). Opponents of this view, call them “No Transfer” theorists, deny that B1 ever transfers to B2. The latter is a rather stringent position. For if B1 never transfers to B2, we are blameless for *all* our ignorant behaviour.⁷ If ignorant agents are blameworthy, No Transfer theorists maintain, they are only blameworthy for their benighting acts (which may, of course, lead to unwitting acts).

⁴ Cf. Wieland (2016). For a critical take on the necessary condition, cf. Clarke (2016), Talbert (2016), Alvarez & Littlejohn (2016).

⁵ This point is due to Holly Smith.

⁶ This explanatory relation is only partial, since B2 has further explanatory conditions (e.g. B2 obtains because she was free to prescribe something else).

⁷ According to Smith (1983, p. 546), if S is blameworthy for an unwitting act it is never because of some B1, but because at the time of the unwitting act there was some better act of inquiry that S could have performed. Smith refers to such actions as “precipitate actions” (2016).

Holly Smith has provided challenging support for No Transfer in her influential (1983) paper and in her (2016) follow-up.⁸ A central part of her defense of No Transfer is that agents who act from ignorance fail to act from insufficient concern. She maintains that acting from insufficient concern is a necessary condition of blameworthiness, and that this condition is not met in unwitting act cases, even when the agent's ignorance traces to a benighting act for which the agent is blameworthy.⁹ In this chapter, we will set this argument aside, and revisit a second argument that Smith considers, namely the argument from *moral luck*.¹⁰ Our aim in this chapter is *not* directly to refute Smith's No Transfer, but rather to determine its strongest opponent. We will show that there are several candidate conceptions of blame transfer and that problems that hold for one conception do not carry over to others.

The plan of the chapter is as follows. In §2, we explain the argument from moral luck. In §3, we introduce the scope/degree distinction in order to show that this argument fails if we assume that transfer involves only an increase in the scope of blameworthiness (rather than also its degree). In §4, we argue that this assumption yields a puzzle. In §5, we explore ways for Transfer proponents to respond to this puzzle, and in §6 we discuss a problem for one of these responses (which relates to the explanatory condition on transfer). In §7, we conclude with an overview of the four accounts of blame transfer that we have explored in this chapter.

2. The argument from luck

Smith's (1983) argument from moral luck for No Transfer rests on the premise that it is a matter of luck whether or not an unwitting act results from a benighting act. Hence, it is a matter of luck whether one is doubly blameworthy (i.e. blameworthy for both the unwitting and the benighting act). *If* this latter claim is unacceptable, then we are only

⁸ Smith (1983, pp. 548-50) uses a three-fold classification. No Transfer corresponds to what Smith calls the "Liberal" view, and Transfer to the "Conservative" and "Moderate" views (which differ regarding issues of degree).

⁹ This assumes that there is no other lack of concern attending the unwitting act apart from the lack of concern directly related to her becoming ignorant. No Transfer proponents would not deny that S is blameworthy for a certain unwitting act if she, for example, falsely believed that the drug would cause some mild harm and prescribed it because she really disliked the patient.

¹⁰ We discuss Smith's concern-based argument for No Transfer in a companion paper "A Puzzle Concerning Blame Transfer."

blameworthy for benighting acts and blameless for any resulting unwitting ones. Here is the argument in steps:

- (P1) For all unwitting acts A, S performs A as a matter of luck.
- (P2) For all unwitting A, if S performs A as a matter of luck, then if S is blameworthy for A, S's degree of blameworthiness is subject to luck.
- (C1) Hence: for all unwitting A, if S is blameworthy for A, then S's degree of blameworthiness is subject to luck. [from P1, P2]
- (C2) Hence: if S's degree of blameworthiness is never subject to luck, then S is blameless for all unwitting A. [from C1]

This argument is valid and meant to constitute a conditional argument for No Transfer. If our degree of blameworthiness is insensitive to luck, then blame never transfers, and we are blameless for any unwitting act. Although Smith (1983, p. 570) mentions that she is tempted to accept the antecedent of this conditional, she does not want to defend this. In the following, we will clarify the two premises, explain Smith's recent qualification of this argument, and finally discuss a challenge to (P2). Though we agree that if S is blameworthy for unwitting acts, then in a certain sense S is subject to luck, we will argue that, this need not imply that S's *degree* of blameworthiness increases and is subject to luck.

Here's the rationale for (P1). In *Lazy Doctor*, the outcome was that the patient died. But the story might well have ended differently. Suppose Julie prescribes *Inscientium*, yet the patient does not have the rare kidney condition and hence does not die. Or suppose Julie prescribes *Inscientium* to a patient with the kidney condition who then forgets to take the drug and does not die. Finally, suppose Julie, through some fluke of scheduling, sees no more hay fever patients, and so never has the chance unwittingly to prescribe *Inscientium* to patients for whom it is deadly. Let us call such cases *Lucky Doctor* cases.

Generally, the occurrence of any wrongful unwitting act (such as prescribing a harmful drug) or bad consequence (such as the patient's death) is contingent and might not have occurred even given the occurrence of the benighting act. This implies that there could be a similar agent (with a similar psychological profile and motives) who performs the same wrongful benighting act, but does not perform any wrongful unwitting act. There could also be a similar agent who performs both the wrongful benighting and

unwitting acts but this does not result in any bad consequence. All such cases count as moral luck cases in the present discussion.¹¹

Smith does not assume further restrictions on moral luck. In particular, she does not assume that cases count as moral luck cases only if most (or a significant portion of) similar agents who perform the same wrongful benighting act either fail to perform the wrongful unwitting act or fail to cause the bad consequence.¹² Suppose the story of Lazy Doctor ended as it did, but that it ended differently for other doctors similar to Julie. Suppose that these doctors also failed to read the article and prescribed Inscientium and that the government intervened just in time preventing the patients' deaths. Comparing Julie's case with the government intervention cases shows how she is subject to luck, but the occurrence of these other cases is not necessary. It suffices for Lazy Doctor to be a case of moral luck if there could be cases where the bad consequence failed to result from her actions.

The rationale for (P2) comes into view by comparing Lazy Doctor and Lucky Doctor. The features of the two cases are identical except that Lucky Doctor's unwitting act fails to result in any deaths. The question here is who is more blameworthy? According to Smith's No Transfer view, Lazy Doctor and Lucky Doctor are blameworthy to the same degree. Lucky Doctor is blameworthy to a certain degree for the benighting act, and Lazy Doctor is no more blameworthy (since the unwitting act does not have any impact on her blameworthiness). But, according to Smith, Transfer proponents would maintain that Lazy Doctor is *more* blameworthy than Lucky Doctor. Lucky Doctor is blameworthy to a certain degree for the benighting act, but Lazy Doctor is blameworthy to a higher degree for the benighting and unwitting act taken together.

Recently, Smith (2016) has revisited this discussion. She no longer thinks that it can be used to resolve the Transfer/No Transfer dispute. Instead, her suggestion is that both views can either deny or accept moral luck, where accepting moral luck is understood as the view that "moral luck in the outcome affects an agent's blameworthiness for performing a risky action" (2016). As Smith puts it, accepting moral luck amounts to the view that consequences can make an agent more blameworthy. In her terms, she gets an extra "black mark" for the benighting act if the subsequent

¹¹ Also known as consequential, resultant or outcome luck.

¹² These similar agents may be actual, or merely possible. For the suggestion that luck has such a modal dimension, cf. Pritchard (2005, ch. 5), Levy (2011, ch. 2).

unwitting act has a bad consequence.¹³ So understood, Transfer proponents and No Transfer proponents can either accept or reject moral luck.

To see how this works, compare Lazy Doctor and Lucky Doctor. In both cases, Julie takes the risk of failing to read the article (at t_1), and then she prescribes the drug (at t_2). The only difference is that for some reason or other the patient does not die in Lucky Doctor. According to Smith, these cases can be analysed in four ways. (a) According to No Transfer proponents who *deny* moral luck, both Julies get a black mark at t_1 , and no black mark at t_2 . (b) According to No Transfer proponents who *accept* moral luck, Julie in Lazy Doctor gets an additional black mark at t_1 on account of the occurrence of a bad consequence. For neither proponent of No Transfer, does Julie receive an additional black mark at t_2 . (c) According to Transfer proponents who *deny* moral luck, both Julies get a black mark at t_1 and one at t_2 . (d) According to Transfer proponents who *accept* moral luck, Julie in Lazy Doctor gets an additional black mark at t_1 , again on account of the bad consequence in her case.

In Smith (1983) and (2016) the assumption is that moral luck (in the sense just discussed) entails an increase in the degree of blameworthiness. Next, we will argue on the basis of the degree/scope distinction that this can be challenged. After that, we will argue that this strategy entails a puzzle, which opens up conceptual space for distinguishing quite different accounts of blame transfer.

3. Scope vs. degree

In Lazy Doctor, a patient dies as a result of Julie's ignorance, whereas no such outcome obtains in Lucky Doctor. The worry is that it is unfair for Julie to be blameworthy for killing her patient in Lazy Doctor, given that in Lucky Doctor (where Julie's psychological profile and benighting act are exactly the same) Julie is not blameworthy for any bad consequence. This is a well-known problem for blameworthiness for outcomes generally. However, it does not immediately render Lazy Doctor blameless for killing the patient.¹⁴ Consider Thomas Nagel's assassin cases (1979). Two assassins try to kill their

¹³ Smith's black marks indicate an increase in degree (or level, as Smith puts it) of blameworthiness, rather than merely its scope. See §3 for more on this distinction.

¹⁴ If you do think that Lazy Doctor is excused for this reason, cf. Khoury (2012), Talbert (2016) and Graham (2016). The basic idea is that outcomes are irrelevant to blameworthiness since they

targets, but one of them fails only because of an interfering bird. Does the fact that the unsuccessful assassin fails to kill render the successful assassin blameless for the killing? It is not at all clear that it does, given that the victim's death was reasonably foreseeable and, we can assume, avoidable.

Similarly, it does not seem to follow from the fact that Lucky Doctor fails to kill any of her patients that Lazy Doctor is blameless for killing her patient, given that prescribing the drug to her hay fever patients was avoidable and that their deaths were reasonably foreseeable. We may safely assume it was avoidable given that there are many different drugs for hay fever on the market (even though they might be less effective). But, was the death of her patient also reasonably foreseeable? As we know since the discussion between Vargas (2005) and Fischer & Tognazzini (2009), the reasonable foreseeability of an outcome depends crucially on its description. The outcome of Julie's failure to read the article admits of at least the following descriptions:

- (1) Julie kills that specific patient by prescribing Inscientium.
- (2) Julie harms an individual at some point in the future by prescribing a drug that was discovered to be dangerous.

Outcome (1) is not reasonably foreseeable in Lazy Doctor given that Julie did not yet read the study and, thus, lacked the information that she would need in order to foresee that Inscientium is fatal for certain people (let alone for that specific patient). Since (2) does not specify the kind of harm, this outcome is reasonably foreseeable. Julie suspects that the article contains information about the safety of Inscientium, a drug relevant to her practice, but she still decides not to read it. Moreover, not keeping up on practice-relevant research is seriously risky in her profession (as the case description is intended to convey).

Of course, not all possible outcomes of her failure to read the article are reasonably foreseeable (under *any* description). Suppose Julie prescribes Inscientium to a patient who on the way to the drugstore causes an accident and kills a pedestrian. Killing a pedestrian on the way to a drugstore is not a foreseeably upshot of prescribing a drug. Still it is an outcome of Julie's conduct: if she had not prescribed the drug, the patient would not have gone to the drugstore.

do not make any difference to the agent's psychology, her moral orientation, or her intentions, which are the sole determinants of blameworthiness on their views.

Even so, there are some tricky issues to be addressed. Suppose Julie is blameworthy for not reading the article (B1), and if blame transfers, she is also blameworthy for killing the patient (B2). Smith (1983, pp. 567-70) asks: does she deserve a double penalty and should we blame her twice?¹⁵ The question is whether transfer proponents are committed to the claim that in Lazy Doctor, her degree of blameworthiness increases.

They are not. It is plausible to maintain that transfer of B1 to B2 does not increase the *degree* of blameworthiness, but merely its *scope*. “Scope” indicates the number of items for which one is blameworthy, while “degree” indicates the magnitude of an agent’s blameworthiness for a specific item (or set of items). The credits of this view go to Zimmerman.¹⁶ According to Zimmerman, Nagel’s assassins are blameworthy to the same degree, though not for the same things. The successful assassin is blameworthy for both the attempt to kill and for the killing, while the unsuccessful one is only blameworthy for the attempt.¹⁷ Similarly, then, Lazy Doctor would be blameworthy for more things than Lucky Doctor, though she would not be blameworthy to a greater degree. We call this the “scoping strategy”.¹⁸ While we develop potential problems for this view below, it is worthwhile exploring as a way to embrace Transfer in the face of considerations about luck.

The scoping strategy also has application in interpersonal cases. Consider a variant of Lazy Doctor in which Julie promises to update a fellow doctor Jaap about the article. As in the original case, she fails to read the article and does not update him. When Jaap is back at work after his sickness, he prescribes Inscientium to a patient who suffers a fatal heart attack as a result. Again, one might ask: for how many things is Julie to be blamed? As Smith asks: “Is the doctor in the original case (in which he performs both the benighting act and the unwittingly harmful act) more blameworthy than he is in the second case, in which he performs only a benighting act, and his colleague performs

¹⁵ Smith refers to Aristotle’s remark that if you did wrong in a state of drunkenness, then according to an ancient law you were to be punished not only for your wrongful conduct, but also for getting into that state of drunkenness (*Nicomachean Ethics*, 1113b30).

¹⁶ See Zimmerman (1988, pp. 133-4; 1997, pp. 418-9; 2002, p. 560).

¹⁷ According to Zimmerman (2002, p. 562), the degree of blameworthiness of the two doctors is the same because Lazy Doctor does not have more control over the occurrence of the outcome.

¹⁸ Interestingly, Smith describes a very similar strategy at the end of her (1983). According to this discussion, Liberals who reject moral luck would maintain that Lazy Doctor is blameless for her unwitting act, even though there is still a legitimate sense in which she would be “to blame” for it.

the resultant unwittingly harmful act?" (2013, p. 3) As we have assumed in §1, blame can transfer not only to unwitting acts (such as prescribing the drug) but also to consequences (such as the death of the patient). In both Lazy Doctor and the variant with Jaap, then, the death of the patient can be seen as a consequence of Julie's omission to read the article. If Julie had read the article (and assuming we are dealing with normal doctors), this outcome would not have occurred. In such cases, the scoping strategy would say that Julie is blameworthy for the consequence in both cases, even though this does not affect her degree of blameworthiness.

4. Puzzle

The foregoing suggests that scope/degree distinction yields a promising view on transfer, however, we will now develop an objection to the scoping strategy. The problem stems from the potential for mismatches between the degree of B1 and the degree of B2. In such cases, it will be problematic to maintain that transfer of B1 to B2 merely increases the scope and not the degree of blameworthiness. These mismatches occur due to factors that intuitively have an impact on B2 but seem not to impact B1, given that they occur only *after* the benighting act. We will refer to these factors as "downstream factors." To illustrate such a mismatch, we appeal to the downstream factor of the degree of badness of the outcome, and we contrast this with a factor that seems to play a role only at the time of the benighting act, namely the difficulty of informing oneself (and so of avoiding the unwitting act).¹⁹ Compare the following cases:

Inconsequential Lazy Doctor. Julie heard that a colleague's recent article might report new findings about the drug Inscientium. In the past month she has not read the required amount of practice-relevant research. She has plenty of time to read this article, but instead chooses to have coffee with a colleague. In the study that Julie fails to read, Inscientium is shown to cause minor, temporary itches in one's feet. The next day, Julie prescribes Inscientium to a patient who is then subject to the itches.

¹⁹ Further potential factors (such as the strength of one's belief that the benighting act is wrong) might yield further puzzles.

Consequential Busy Doctor. Julie heard that a colleague's recent article might report new findings about drug Inscientium. In the past month she has not read the required amount of practice-relevant research. She could have skimmed the article in her 30 minute break, but forgets to do it due to her busy schedule. In the study that Julie fails to read, Inscientium is shown to cause heart attacks. The next day, Julie prescribes Inscientium to ten patients who then all die from a heart attack.²⁰

In Inconsequential Lazy Doctor, Julie seems blameworthy to a high degree for not reading the article. After all, she has plenty of time to read this study, and, keeping fixed certain other relevant factors, it seems that the easier it is to avoid the benighting act, the higher one's degree of B1.²¹ Yet, one might wonder whether this high degree of B1 transfers to a high degree of B2 for the patient's itches.

In Consequential Busy Doctor, by contrast, Julie seems blameworthy to a low degree for not reading the article. She is still blameworthy because she has not read the required amount of research. Perhaps Julie could have planned her time better, but in her current circumstances it is difficult for Julie to read the study. Here, one might wonder whether this low degree of B1 transfers to a low degree of B2 for the ten deaths. These cases are especially interesting in comparison. Who is overall more blameworthy, Inconsequential Lazy Doctor or Consequential Busy Doctor?²²

One might think that this question, and particularly the notion of blameworthiness at issue, is underspecified. For our purposes, we will remain mostly non-committal as to which notion of blameworthiness is the operative one in this context. Importantly, though, we are not interested in a consequentialist account, which might entail that we should punish one of the Julies more severely given the favorable outcomes of such punishment (e.g. punishment might give other doctors an incentive to keep up on practice-relevant research, or it might help the families of the victims to deal with their losses.) We are interested more broadly in accounts of blameworthiness that

²⁰ The consequences in each case (i.e. the itches and heart attacks) fall within the range of what was reasonably foreseeable depending on their description.

²¹ For more on difficulty, see Bradford (2016) and Guerrero (2016).

²² "Overall blameworthiness" is blameworthiness for a set of items. We do not assume that the degree of overall blameworthiness for a set is something over and above the sum of the degree of blameworthiness for its members. Hence the overall degree of B1 and B1 taken together ('B1+B2') is composed of the degree of B1 plus the degree of B2.

understand the question as which doctor deserves a harsher reaction given what happened as a consequence of her conduct. Who deserves to be resented more and by more people? According to a Scanlonian picture, for example, the outcomes of faulty actions can increase the significance of the fault and hence the degree of blameworthiness by increasing the number of moral patients to whom it is appropriate for the agent to make amends (Scanlon 2008, p. 150).

Given the scoping strategy explained in §3, one might think that the degree of B1 determines one's overall degree of blameworthiness and that B2 only enlarges the scope of blameworthiness. This would imply that Inconsequential Lazy Doctor is overall more blameworthy than Consequential Busy Doctor (i.e. that the degree of B1+B2 in the former case is higher than the degree of B1+B2 in the latter.) One might react in two ways to this implication:

1. Yes, Inconsequential Lazy Doctor is more blameworthy overall than Consequential Busy Doctor.
2. No, Consequential Busy Doctor is more blameworthy overall than Inconsequential Lazy Doctor.

Both responses are plausible. On the one hand, Inconsequential Lazy Doctor seems more blameworthy given her laziness. On the other hand, Consequential Busy Doctor seems more blameworthy given the all the upshots of her behaviour. The first response is compatible with the scoping strategy that assumes that the degree of badness of the outcome (namely, ten deaths, as opposed to temporary itches) is irrelevant for blameworthiness. The degree of B1 plus a potential increase in scope is all that is needed to understand transfer. The second response does not seem compatible with the scoping strategy. For it assumes that the degree of badness of the outcome does matter, and that one's overall degree of blameworthiness is not only dependent on how much time one had to inform oneself (and other such factors at or before the benighting act), but also on how bad it turned out to be that one failed to do so.

5. Alternative views

Of course allowing for the badness of outcome to influence blameworthiness presupposes a particular position on the problem of luck, namely that outcomes can affect one's overall blameworthiness. Like Smith (1983, pp. 567-70), we will not take

sides in this major debate. Our goal is to show that the second response (i.e. that Consequential Busy Doctor is more blameworthy) might in fact be spelled out in two very different ways. If Consequential Busy Doctor is indeed blameworthy to a greater degree, then the question arises where the high degree comes from.

There seem to be at least two options. First, one may accept that transfer of B1 to B2 increases only the scope of blameworthiness, and not the degree, but add the claim that the degree of B1 (and so of B2) is partly dependent on the badness of the given outcome. On this proposal, Consequential Busy Doctor is not only blameworthy to a greater degree for the outcome, but also for not reading the article.²³

We think this option is problematic for two reasons. First, consider that Consequential Busy Doctor was really busy. The amount of effort it would have taken her to read the article seems to be a *prima facie* mitigating factor. If this is right, then it is simply mistaken to claim that her degree of B1 is high. Why think that the outcome's contribution to the degree of B1 (which, according to this option, would be really high, because of the deaths) swamps the "agent's contribution" to the degree of B1 (which would be really low, because she was busy)?²⁴

Second, how could the badness of the outcome of one's unwitting act possibly affect the degree of B1? When Julie omits to read the article at t_1 , there either is a fact of the matter about the degree of B1 or there is not. If there is not, then the degree of B1 is only established at t_2 , when the outcome occurs. But this seems strange. For even if Julie drops dead immediately after having the coffee with her colleague, it seems that we can say she is blameworthy to some degree or other. Notice that this counterintuitive result is absent in normal "backwards effect" cases. Say I run a marathon in four hours on Monday, and that you run one in three hours on Tuesday. On Monday there is no fact of the matter about whether I ran a slower marathon than you did, but on Tuesday there such a fact of the matter, and so it is only on Tuesday that my running a marathon in three hours on Monday becomes my running a slower marathon than you did. Since there is nothing controversial about there being on Monday no fact of the matter about

²³ To a certain extent, this strategy is similar to what Smith (2016) sees as the acceptance of moral luck, which we discussed in §2 above.

²⁴ A proponent of this view might respond to this by claiming that the outcome's impact on degree of B1 is *constrained* in some way by the degree of B1 that would obtain even if the agent never acted on the resulting ignorance. In fact, we consider think such a strategy promising (as we will discuss later), though we will not explore it here since we think the view suffers from a second worry.

whether I ran a slower marathon than you, given that you had not yet run yours, one might think that it is similarly uncontroversial that at the time of the benighting act, there is no fact of the matter about whether Julie is blameworthy for failing to read the article given that she has not yet performed any subsequent ignorant action. The crucial difference between these two cases, of course, is that it is implausible to think that Julie is not already blameworthy for failing to read the article during the interval of time between this failure and her unwitting killing, while it is not at all implausible to think that I have not run a slower marathon than you during the interval of time between my running and your running. This shows that the view under consideration entails a different and potentially more problematic sort of backwards effect than what is found in the marathon case.

As an alternative explanation for how badness of outcomes can increase degree of blameworthiness for the benighting act, we may say that at t_1 there *is* a fact of the matter about the degree to which Julie is blameworthy for her benighting act—perhaps a high degree in the lazy case, but a low degree in the busy case—but then at t_2 , once the unwitting act is performed and its bad outcome realized, the degree of blameworthiness of B1 increases. On this view, when the outcome of the unwitting act occurs, some blameworthiness fact about the past changes. But, this also seems strange. This case is again distinct from ordinary backwards effect cases, where no facts about the past change with the occurrence of later events. It is still the case on Monday that I ran a marathon in four hours even if on Tuesday it becomes the case that on Monday I ran a slower marathon than you did. On the view under consideration, however, Julie's degree of B1 has a low value at t_1 and then at t_2 , this fact about her changes—she becomes blameworthy to a higher degree. The oddness of this implication is brought out by the fact that someone who only mildly blamed Consequential Julie at the time of their benighting act, perhaps by feeling a only a modicum of vicarious resentment, would at the time of the outcome of Julie's very consequential unwitting act realize that the intensity of her indignation actually failed to match how blameworthy Julie actually was at the prior time of the benighting act. But, it is plausible to suppose that this person in this case made no mistake in her appraisal. And, note that this problem cannot be avoided simply by saying that at the time of the outcome, the appraiser is justified in ramping up the intensity of her indignation to match the increase in degree of blameworthiness. This may be an appropriate reaction in the light of the entire sequence—from benighting act to the consequential outcome—but, on

the view under consideration, there would still have been a mismatch between the high degree of B1 and the appraiser's reaction at the time of the benighting act.²⁵

If these objections are on the right track, then there are reasons not to take a modified scoping strategy according to which the degree of B1 that scopes can be affected by downstream factors like badness of outcome. A second way of responding to Inconsequential Lazy Doctor and Consequential Busy Doctor is to reject the scoping strategy and accept that transfer of B1 to B2 *can* increase the degree of overall blameworthiness when the badness of the outcome is of sufficient magnitude. On this proposal, the degree of blameworthiness for the benighting act is not affected by the fact that the unwitting act occurs, but the degree of blameworthiness for the unwitting act does depend on the badness of the outcome. Thus, Inconsequential Lazy Doctor and Consequential Busy Doctor are blameworthy to a certain degree of their benighting act (depending on factors obtaining at or before this act), though Consequential Busy Doctor is blameworthy to a greater degree for her unwitting act (given that ten patients died as a result of her unwitting act.) Since on this view, only the degree of B2, and not the degree of B1, is affected by such downstream factors, this approach lacks the problematic features of the modified scoping strategy.

On first sight, Consequential Busy Doctor may cast doubt on the idea of blame transfer in general, since “the smallish degree of blameworthiness that is ‘transferred’ cannot explain the large degree of blameworthiness that accrues for the unwitting act”.²⁶ Yet, it is important to see that this worry only applies to the original scoping strategy. The alternative views that we have developed in this section can account for the difference in degree, and their rationale is in part that they are able to do this. Also, it is noteworthy that No Transfer does not fare any better here. For according to No Transfer, agents are only blameworthy for benighting acts, and as such, Consequential Busy Doctor does not deserve much blame at all.²⁷

²⁵ It is true that we often realize that we were not sufficiently horrified by some event (e.g. a campus shooting) until all the facts have come in. But in such cases the mismatch is often grounded on the difference between how horrible we believed something was at the time of our initial reaction and how horrible it actually was. In this case, the mismatch is grounded on the difference between how blameworthy we knew someone to be for her benighting act at t_1 (low) and how blameworthy we knew them to be for that same act at t_2 (high).

²⁶ We are thankful to Holly Smith for this point.

²⁷ This depends, of course, on whether No Transfer accepts moral luck (see §2). If it does, then Consequential Busy Doctor might be quite blameworthy indeed for the benighting act. This

6. Swamping problem

Although the non-scoping view just discussed has an advantage over the modified scoping view, it suffers from what we call the ‘swamping problem’. Generally put, if the magnitude of the outcome’s badness is great enough, the degree of B2 can ‘swamp’ the degree of B1. In such a case, the contribution of B1 to an agent’s overall degree of blameworthiness is negligible. Consequential Busy Doctor already possesses this feature to some degree, but a more extreme case reveals this problem more clearly:

Extremely Consequential, Extremely Busy Doctor. Julie heard that a colleague’s recent article might report new findings about drug Inscientium. In the past month, in addition to being busy with work, she had to manage a very pressing and emotionally exhausting situation at home, and she recently discovered that she will likely be laid off, which meant that she had to begin immediately a time-consuming job search. Despite all of this, she found the time to read almost all of the required amount of practice-relevant research, though she is still 30 minutes short of the required target. She could have skimmed the article about Inscientium in her 30 minute break, but she forgets to do it due to her busy schedule and preoccupied thoughts about her family and employment situation. In the study that Julie fails to read, Inscientium is shown to cause heart attacks. Over the next two days, Julie prescribes Inscientium to 30 patients who all die from a heart attack several days later.

In this case, Julie is in a really bad spot. The various things going on in her life make it difficult for her to comply with her professional obligation to stay current on her research. Still, she overcomes all of these factors and very nearly complies with it. But, near compliance is still short of compliance, and we can stipulate that the difficulty of compliance does not defeat the presence of the obligation. Certain moral obligations are just that demanding. Given Julie’s circumstances however, it is plausible to suppose that

version of No Transfer would be vulnerable to the same objection as the first view explored in this section.

the degree of B1 is maximally low or else close to it.²⁸ As before, this is especially clear in comparison with the various Lazy Doctor cases above, wherein these blameworthiness mitigating factors are completely absent. What about B2, then? On the assumption that badness of outcome is a B2 aggravating factor, the deaths of 30 patients would result in a very high degree of B2. Julie would be very blameworthy indeed given that so many people died as a result of her actions. It follows then that Julie's overall blameworthiness (B1+B2) is very high, and this is despite the fact that her degree of B1, which is, so to speak, the original sin, is maximally low. In such a case, B1 is swamped by the degree of B2, meaning that when one considers overall blameworthiness for the relevant sequence of events, B1's contribution is negligible. We think this is a problem because if one requires that B1 is an explanatory condition for B2 (as is plausible given what we say in §1), then it should matter whether or not B1 is fulfilled. But if B1 is swamped by B2, it makes little sense to say that B1 explains B2.

One way of avoiding the swamping problem is to maintain that while B2 can have an impact on overall degree of blameworthiness, mediated by factors such as badness of the outcome, its impact is *limited* by the degree of B1. There are a number of possible ways in which the limitation might work. For example, the degree of B1 might set the upper bound of degree of B2 or, alternatively, some multiple of the degree of B1 might set the upper bound of B2. Since it is beyond the scope of this chapter to defend one of these proposals over the other, we adopt the former proposal in what follows. To see how this works, it helps to translate degrees of blameworthiness to some quantifiable metric.²⁹ We stipulate that the degree of blameworthiness for a given agent's action is a function of the number of "demerits earned" by the agent. On this view, Extremely Consequential, Extremely Busy Doctor earns only two demerits for failing to read the article—her degree of B1 is very low. This degree of B1 constrains the number of demerits Julie can receive for her unwitting act (and so the degree of B2). Thus, at most, she can earn two more demerits for B2, which would give her a total of four demerits.

Note that this is distinct from the scoping strategy since it makes a difference to Julie's overall degree of blameworthiness if she performs the unwitting act and thereby

²⁸ If this case is not convincing, replace it with a case that has the following profile: factors attending the benighting act make for extremely low degree of B1 and factors attending subsequent unwitting act make for extremely high degree of B2.

²⁹ This is overly precise, and perhaps comically so, but nothing hangs on there being this degree of precision. The same account can be developed using some other function for determining degrees of blameworthiness, such as the intensity of some fitting reactive attitude.

realizes the relevant bad outcomes. If something would have intervened, Julie's patients would not have died, and she would not earn any additional demerits—her overall number of demerits would be lower (two, say, instead of four). But, on this view, the degree of B1 cannot be swamped by degree of B2, and thus it retains its explanatory significance to the issue of overall blameworthiness. This account preserves the importance of B1 in determining degree of overall blameworthiness, while still leaving room for downstream factors to affect B2. Indeed, one of the most important of these factors is the degree of B1, which acts as a constraint on how much B2 adds to overall blameworthiness.

Let us now apply this machinery to our two puzzle cases. In Inconsequential Lazy Doctor, Julie's failure to read the article was really gratuitous. This means that we can suppose she receives five demerits for it. This complies with the intuition that B1 is greater in her case than in any of the Busy Doctor cases. It also sets the upper bound for B2, so that if something really consequential were to happen she would receive five more demerits for B2 for a total of ten. In this case, though, something rather inconsequential happens, namely her patients suffer mild, temporary foot-itchiness. Thus, we can say she receives only one demerit for her unwitting act, which yields a total of six demerits. In Consequential Busy Doctor, by contrast, Julie's failure to read the article was slightly less gratuitous, so we can say she receives four demerits for it, and hence a maximum of eight demerits in total. Given that ten of her patients die, she does receive all eight demerits. According to this system, then, Inconsequential Lazy Doctor is less blameworthy (six demerits) than Consequential Busy Doctor (eight demerits).³⁰

Hence, there are actually two versions of the view that B2 contributes independently to overall blameworthiness. On the first version, the degrees of B1 and B2 can be determined independently and simply summed to determine overall blameworthiness. On the second version, these degrees of blameworthiness are also summed, but the degree of B1 constrains the degree of B2, which makes it resistant to the swamping problem.

7. Conclusion

³⁰ To be sure, different metric systems will yield different verdicts. If Consequential Busy Doctor receives only two demerits for the benighting act, she will receive maximally four demerits for her overall conduct (and so on).

If blame transfers in certain cases, then what does that mean? In this chapter, we have multiplied the options for Transfer proponents. In particular, we have distinguished four views:

- (I) The degree of B1 determines the overall degree of blameworthiness; transfer to B2 only enlarges the scope of blameworthiness, and the degree of B1 is insensitive to downstream factors.
- (II) The degree of B1 determines the overall degree of blameworthiness; transfer to B2 only enlarges the scope of blameworthiness, yet the degree of B1 (and so of B2) is sensitive to downstream factors.
- (III) The degree of B1 and of B2 determines the overall degree of blameworthiness; the degree of B1 is insensitive to downstream factors; the degree of B2 is sensitive to downstream factors and is not constrained by the degree of B1.
- (IV) The degree of B1 and of B2 determines the overall degree of blameworthiness; the degree of B1 is insensitive to downstream factors; the degree of B2 is sensitive to downstream factors, though it is constrained by the degree of B1.

We think that characterizing these possible Transfer accounts is a substantial result, even we leave certain main issues unaddressed. For one thing, nothing we have argued here decisively supports one view or other. In §§5-6, we argued that (II) and (III) have some counterintuitive consequences, so we think the real contenders are (I) and (IV). In addition, Smith's No Transfer position is still a live option (even though it does not follow from the argument from luck, as Smith acknowledges.) Additionally, even assuming we accept Transfer, the question of whether (I) or (IV) is the better view will depend in part on the relevance of downstream factors to blameworthiness, and this is an issue that we have not tried to settle in this chapter.

In closing, however, let us just mention one independent consideration that might be invoked to settle the debate between (I) and (IV). One might think that if Julie is only derivatively blameworthy for prescribing the drug and killing the patient (because her blameworthiness derives from her blameworthiness for not reading the article), then her blameworthiness would reduce to an "intellectual crime".³¹ In such a case, we would think of Julie as someone whose main fault is that she did not inform herself. She is at bottom someone who failed to look into something, and her degree of blameworthiness is calibrated to that which is appropriate for this kind of intellectual failure. But, one might

³¹ This point is due to Liz Harman.

find this implausible. In Julie's case people died as a result of her unwitting conduct. For this reason, we may think of Julie as someone whose main fault is that she failed to prevent avoidable deaths and failed generally to care about the potentially bad consequences of her ignorance. Her degree of blameworthiness should be calibrated to that which is appropriate for this kind of severe moral failure.

If one thinks that Julie commits what is really an intellectual crime, then this is consistent with (I) which states that the degree of blameworthiness for one's unwitting conduct is fully determined by the degree of blameworthiness for not informing oneself. Here, one might say that, in a relevant sense, B2 is a mere intellectual crime. If one thinks instead that Julie's unwitting act is not a mere intellectual crime, but something more grave, then this is only consistent with (IV), which holds that the degree of B2 is not only dependent on the degree of B1, but also on downstream factors. While we do not take these breezy comments to be at all decisive in favor of (IV), we hope they provoke further inquiry into the relevance of downstream factors to questions of blameworthiness.

Finally, certain readers might not be attracted to the overall B1/B2 picture in the first place. In particular, King (2016) argues against the tracing/transfer way of framing things. He argues that all unwitting agents are either reckless or negligent and that these concepts suffice for an analysis of blameworthiness. King's basic strategy is to diffuse issues of scope and to argue that blameworthiness for unwitting acts and blameworthiness for benighting acts are no longer treated as separate questions. There is just *one* question of overall blameworthiness. Lazy Doctor is overall blameworthy for her reckless conduct.

In response, we think the B1/B2 picture has certain advantages. Cases such as Inconsequential Lazy Doctor and Consequential Busy Doctor show that to a certain extent blameworthiness for benighting acts and blameworthiness for unwitting acts (and their consequences) might well come apart. They also show that the question of overall blameworthiness poses a difficult puzzle. The B1/B2 picture is helpful here because it allows us to see this.

One might even take a more radical position and hold that comparative questions regarding overall blameworthiness do not make sense in the first place (or are at least less fundamental than questions regarding B1 or B2.) Indeed, no puzzle arises if we just ask the two separate questions about B1 and B2 and compare benighting acts of a certain kind ('Lazy Doctor is more blameworthy for not reading the article than Extremely Busy Doctor') or unwitting acts of a certain kind ('Consequential Doctor is more blameworthy for prescribing the drug than Inconsequential Doctor'). In our view,

though, this move should be resisted. As all transfer cases show, we cannot answer B2 questions (and determine whether Consequential Doctor is more blameworthy, if she is blameworthy at all, for prescribing the drug than Inconsequential Doctor) without considering the answer to B1 questions, namely the degree to which agents are blameworthy for their ignorance. This is exactly where our views on blame transfer become relevant.³²

References

- Alvarez, M. & C. Littlejohn 2016. When Ignorance is No Excuse. This volume.
- Bradford, G. 2016. Hard to Know. This volume.
- Clarke, R. 2016. Ignorance, Revision, and Commonsense. This volume.
- Fischer, J. M. & M. Ravizza 1998. *Responsibility and Control. A Theory of Moral Responsibility*. CUP.
- Fischer, J. M. & N. Tognazzini 2009. The Truth about Tracing. *Noûs* 43: 531-56.
- Guerrero, A. 2016. Intellectual Difficulty and Moral Responsibility. This volume.
- Graham, P. 2016. The Epistemic Condition on Moral Blameworthiness, A Theoretical Epiphenomenon. This volume.
- Khoury, A. C. 2012. Responsibility, Tracing, and Consequences. *Canadian Journal of Philosophy* 42: 187-207.
- King, M. 2016. Tracing the Epistemic Condition. This volume.
- Levy, N. 2011. *Hard Luck. How Luck Undermines Free Will and Moral Responsibility*. OUP.
- Nagel, T. 1979. *Mortal Questions*. CUP.
- Pritchard, D. 2005. *Epistemic Luck*. OUP.
- Robichaud, P. & J. W. Wieland 2017. A Puzzle Concerning Blame Transfer. *Philosophy & Phenomenological Research*.
- Rosen, G. 2004. Skepticism about Moral Responsibility. *Philosophical Perspectives* 18: 295-313.
- Scanlon, T. M. 2008. *Moral Dimensions. Permissibility, Meaning, Blame*. HUP.
- Smith, H. M. 1983. Culpable Ignorance. *Philosophical Review* 92: 543-71.
- 2013. Negligence. In *The International Encyclopedia of Ethics*. Wiley.

³² We would like to thank Michael Zimmerman and especially Holly Smith for very helpful comments.

- 2016. Tracing Cases of Culpable Ignorance. In R. Peels ed., *Moral and Social Perspectives on Ignorance*. Routledge.
- Talbert, M. 2016. Akrasia, Awareness, and Blameworthiness. This volume.
- Vargas, M. 2005. The Trouble with Tracing. *Midwest Studies in Philosophy* 29: 269-91.
- Wieland, J. W. 2016. The Epistemic Condition. This volume.
- Zimmerman, M. J. 1988. *An Essay on Moral Responsibility*. Rowman & Littlefield.
- 1997. Moral Responsibility and Ignorance. *Ethics* 107: 410-26.
- 2002. Taking Luck Seriously. *Journal of Philosophy* 99: 553-76.