

The Impermissibility Criterion on Blameworthiness

Abstract: In this paper I defend the view that behaving impermissibly is necessary ground of being blameworthy (“the impermissibility thesis”). First, I offer a brief positive argument for this thesis: that we ought to believe it because it is the principled explanation of our intuitive judgments about blameworthiness. I spend the remainder of the paper defending the impermissibility thesis against an objection from T.M. Scanlon. Scanlon thinks that agents can be blameworthy without behaving impermissibly in cases where they “do the right thing but for the wrong reason.” Only what is under an agent’s voluntary control can be impermissible, Scanlon thinks, and yet agents cannot choose the reasons for which they perform a given action. He believes, therefore, that the reasons for which an agent performs a given action can sometimes make that agent blameworthy without making her action impermissible. I argue that Scanlon is correct that only what is under an agent’s voluntary control can be impermissible. He is also correct that there is a sense in which agents cannot choose the reasons for which they perform a given action: they cannot choose their motivating reasons by *choosing to be motivated by some considerations rather than others*. Nevertheless, although agents cannot choose to be motivated by some considerations rather than others, they can voluntarily control which considerations they are motivated by via their choices about *what to do*. Therefore, the reasons for which an agent performs a given action *do* affect the permissibility of that action even when Scanlon thinks they do not. Accordingly, cases in which agents are blameworthy for acting for particular reasons can be read as cases in which agents act impermissibly, and do not demonstrate that impermissibility and blameworthiness come apart.

I. The Impermissibility Thesis

People are blameworthy only when and because they have committed some wrongdoing for which they are responsible. In other words, doing something impermissible, for which one is responsible, is a necessary ground of being blameworthy. Most people accept at least part of this claim. It is widely believed – and from now on I will assume – that agents are only blameworthy for what they’re responsible for. Let’s call this the responsibility criterion on blameworthiness. Here, I want to defend an impermissibility criterion on blameworthiness: that agents are blameworthy only when and because they have behaved impermissibly. I’ll refer to this claim as the “impermissibility thesis.”

The impermissibility thesis may strike one as obviously true. Indeed, that is how it strikes me. The basic idea, after all, is that people can’t be blameworthy unless they’ve done

something *wrong*.¹ But in spite of its intuitive appeal, the impermissibility thesis requires defense, because philosophers have recently offered arguments to the contrary.² A main goal of this paper will be to respond to just such an argument from T.M. Scanlon.

Scanlon's proffered reason for rejecting the impermissibility thesis is that he thinks that agents can be blameworthy for "doing the right thing but for the wrong reason."³ Specifically, he believes that only what is under an agent's voluntary control can be impermissible, and yet that the reasons for which an agent acts are sometimes not under her voluntary control. Thus, he thinks that there are cases in which an agent's reasons for acting make her blameworthy without affecting the permissibility of her behavior. My strategy for rebutting his argument will involve explaining how an agent's motivating reasons *are* under her voluntary control, and hence how they *are* relevant to the permissibility of what she does, even in cases where Scanlon thinks they are not. I hope that this discussion will be of interest not merely for the role it plays in defense of the impermissibility thesis, but also simply because it investigates the type of control we have over our reasons for action.

But, before I come to defending the impermissibility thesis against Scanlon's refutation of it, I'd like to formalize the positive case for its acceptance. So here's the plan: In section II, I'll offer a brief positive argument for the impermissibility thesis: that we ought to accept it because it is the principled explanation of our intuitive judgments about blameworthiness.

¹ Agents do something "impermissible" in the sense I intend when they violate practical obligations to which they are subject. Thus, failing to believe what the evidence suggests isn't "impermissible" in my sense. However, I'll leave open the questions of *what* practical obligations people have, and of whether those obligations are always *moral*.

² Scanlon (2008; 2013) and Capes (2012). Wolf (2011, 332) endorses Scanlon's position on blameworthiness and impermissibility, and Smith (2013) adopts Scanlon's account of blameworthiness.

³ (2008), 57

Then, in section III, I'll present Scanlon's argument against the impermissibility thesis, and, in section IV, I'll reply to it. In section V, I'll offer some concluding remarks.

In defending the impermissibility thesis, I don't want to claim that everything agents can be blameworthy *for* must be impermissible, since agents can be blameworthy for all sorts of things, some of which – like getting in a car accident – it would be strange to call impermissible. However, if what an agent is blameworthy *for* is not itself impermissible, then I think she is only blameworthy for it because it can be “traced” back to something impermissible she has done – like driving recklessly. Hence the technical version the thesis I'll defend is: An agent is blameworthy for some x only if, and because, x is either an impermissible action or omission, or the consequence of one. Throughout this paper I will use the term “object” to refer to those things that agents can be blameworthy *for*. It seems to me that agents can be blameworthy for all sorts of objects including, but not necessarily limited to: actions, omissions, events, attitudes, feelings, lacking attitudes, lacking feelings, states of affairs, and perhaps even the actions of others in their care.

II. The Intuitive View

We should believe that behaving impermissibly is a necessary ground of blameworthiness because that thesis is the principled explanation of our intuitions about blameworthiness: our intuitions about whether agents are blameworthy track whether those agents have done anything impermissible. This is true even in cases where agents are potentially blameworthy not directly for what they have done, but rather indirectly for the consequences of their behavior.

Consider, for instance, whether I would be blameworthy for accidentally spilling a glass of red wine on your carpet. To see how our blameworthiness intuitions track impermissibility, we must fix that we are only considering cases in which I am responsible for the spill, which we can do so by stipulating that the spill was a *reasonably foreseeable* consequence of some earlier commonplace voluntary action of mine.⁴ (If we don't hold my responsibility fixed, then we might get a result in which, although I act impermissibly, I am not blameworthy – but not because blameworthiness *isn't* tracking impermissibility, but rather because the responsibility criterion that I mentioned earlier hasn't been met.) So let's stipulate that this is not a case in which someone else carried the glass into the living room and placed it just behind my elbow unbeknownst to me. Rather, I voluntarily carried the wine into the living room myself.

Holding fixed that I did take the glass into the living room myself, I still won't be intuitively blameworthy unless my doing so was impermissible. If it was, say because you have an explicit rule against drinks in that room, or because the living room carpet is white and I know that you care about it, then I'll be blameworthy. But if, say, the carpet is dirty and other people are in the living room holding cocktails, such that I had no reason to think you had a preference against drinks in that part of the house, then carrying my glass into the living room will not have been impermissible, and I won't intuitively be blameworthy. That, even when an agent is responsible, she still won't be intuitively blameworthy unless she has

⁴ For arguments that individuals can be indirectly blameworthy for the reasonably foreseeable consequences of their voluntary actions and omissions, see Fischer and Ravizza (1998); Fischer and Tognazzini (2012). Some philosophers would dispute that agents can be responsible for objects that aren't actions or omissions *only* if such objects can be "traced" back to some previous action or omission. (Smith 2005) However, even these philosophers can agree that demonstrating that an object is the reasonably foreseeable consequence of a voluntary action or omission is at least *one way* to demonstrate an agent responsible for it.

also done something impermissible, should make us think that doing something impermissible is a necessary ground of blameworthiness.

Furthermore, we can *predictably* vary our intuitions about whether agents are blameworthy for the results of their behavior by varying whether those objects are the reasonably foreseeable consequences of some impermissible deed. Consider, for instance, a case in which the potential object of blameworthiness is the absence of a feeling: the lack of love for another person.

Generally people are not blameworthy for not loving me. Of course, I may *want* someone to love me, and, depending on who the person is, there are a number of negative reactions I may have upon learning that she doesn't. I may be hurt, upset, or angry, or even dislike a person because of it. However, I can experience all of these attitudes in the absence of blame, and although I may disregard or even scorn a person who doesn't love me, it is generally not the case that she is blameworthy on account of simply not loving me.

There are, however, particular circumstances in which an individual can be blameworthy for not loving me. Usually this will happen only if we have a unique sort of relationship with one another: My mother, perhaps, or my husband can be blameworthy for not loving me. But these are just cases in which, because of our special relationships, agents have special obligations towards me. Hence, it seems likely that people are not normally blameworthy for not loving me because they are not normally under obligations to do so, and that in cases where agents *are* blameworthy for not loving me, they are because their lack of regard indicates a failure to abide by their obligations towards me. If my husband ceases to love me,

for example, I may regard him as blameworthy because I think he has failed to fulfill an obligation of our relationship.

A complication here is that it seems implausible that anyone can be obligated to love me, since loving me is not generally something that a person can choose to do. And in fact, I think that my husband does not have an obligation to love me. However, he does have an obligation to do whatever is in his power to maintain his love for me. When a couple takes marriage vows, they don't undertake obligations to love one another indefinitely. Rather, they undertake obligations to do their best to foster their love for one another – to put themselves in situations conducive to the maintenance of their love. Thus, if my husband is blameworthy for not loving me, it is because his lack of love is indicative of a failure to abide by his obligations to maintain and cultivate his regard for me.

More specifically, my husband can violate his obligation to foster his love for me by performing some action of which his ceasing to love me is a reasonably foreseeable consequence, and, intuitively, he will be blameworthy for not loving me when this is what has happened. So if (in the midst of our divorce) I express blame for my husband's lack of love, and he responds that my blame isn't warranted because he just can't help his feelings, I might reply that in fact it *is* warranted, because if he had taken the job he wanted as I suggested instead of choosing the job he didn't want in order to spend more time at home then he wouldn't be so resentful of me now for his unhappiness with his career. Of course I may be *wrong*. But that is not the point. Rather the point is that I can only see my husband as blameworthy for not loving me, insofar as I see that failure as a reasonably foreseeable consequence of his impermissible deed.

In fact, the possibility that I am wrong brings out a crucial feature of the case: The more convinced I become that my husband has *not* done anything impermissible, the less blameworthy he becomes in my eyes. The more I come to agree that he just couldn't foresee the emotional toll of his job, and therefore that taking it was permissible, the less sustainable my judgment that he is blameworthy becomes. I may still be angry, upset or sad. But I can no longer see him as blameworthy. And vice versa: the more able I am to pinpoint some wrongdoing of my husband's of which the attitude that upsets me was a reasonably foreseeable consequence, the more I see him as blameworthy. Thus, we can actually predictably vary my (and others') judgment about whether my husband is blameworthy for not loving me by varying whether this object was the reasonably foreseeable consequence of some impermissible action or omission of his.

A worry: a subtle feature of this case is that the reasonable foreseeability of my husband's lack of love actually seems to matter twice. That it was the reasonably foreseeable consequence of some earlier action is *both* what seems to make my husband responsible for his lack of love *and* what seems to make his earlier action impermissible. So one might agree that my judgment regarding my husband's blameworthiness for lack of regard predictably varies along with my judgment regarding whether that lack is the reasonably foreseeable consequence of some choice he has made, and yet worry that this is only because my blameworthiness judgment is varying along with whether his lack of regard is something for which he is *responsible* – and not also whether he has done something *impermissible*. But we can confirm that my judgment is, in fact, once again tracking impermissibility by reflecting on variations of the case. For instance, my blameworthiness intuition will dissipate not merely if

I decide that my husband couldn't foresee that taking a job he didn't want would lead to his resenting me, but also if I decide that although he could foresee the risk, his obligations to avoid it were outweighed by other obligations (such as obligations to spend time with his family). So it does seem as if we can predictably vary my blameworthiness judgment specifically by varying whether my husband's lack of regard is the reasonably foreseeable consequence of some *impermissible* deed.

Holding fixed an agent's responsibility, intuitive judgments about that agent's blameworthiness seem to depend on whether she has behaved impermissibly. Furthermore, we can predictably vary blameworthiness intuitions by varying whether a potential object of blameworthiness is the reasonably foreseeable consequence of some impermissible deed. That our intuitions about blameworthiness track our intuitions about impermissibility in these ways, I submit, provides very strong evidence for the thesis that behaving impermissibly is a necessary ground of blameworthiness.

III. Scanlon's Argument against the Impermissibility Thesis

Nevertheless, Scanlon thinks the impermissibility thesis is false, because he thinks it is subject to counterexamples. In particular, Scanlon insists that agents can be blameworthy without acting impermissibly because he thinks this must be the correct description of agents who do "the right thing but for the wrong reason."⁵ He proposes the following example:

Rescue to Riches

A person that I hate, and who I would be happy to see die, is drowning. I save him, but I am able to bring myself to do so only because I do not want him to die right now, since that would mean that his heir, with

⁵ Scanlon (2008), 57

*whom I am currently locked in a bitter electoral contest, would inherit a large sum of money to spend on her campaign.*⁶

It may seem intuitive that I am acting permissibly in this case, but for bad reasons. If a drowning person requires rescuing it can't be *impermissible* to save him – after all, what am I supposed to do, *not* save him? But my reasons for saving him are dreadful. Scanlon's gloss is that, while I act permissibly, I am still blameworthy because my reasons reveal something about my character that impairs my relationships with others. Thus “right thing for the wrong reason” cases show that agents can be blameworthy without behaving impermissibly.

This result depends on reading the case as one in which I *do* act permissibly. And we might instead suggest that my motivating reasons here make my action impermissible: that the permissible action would be to save the drowning person for the *right* reasons. But Scanlon thinks this cannot be correct, because of the relationship between permissibility and choice:

The question of permissibility is the question, ‘May I do X?’ which is typically asked from the point of view of an agent who is presented with a number of different ways of acting. The question is, which of these may one choose? The question of permissibility thus applies only to alternatives between which a competent agent can choose.⁷

Furthermore, he argues, agents cannot choose to perform a given action for one reason rather than another:

When one sees each of several courses of action as supported by sufficient reasons, one can choose which of them to take. In this sense one can choose which reasons to act on. But...when one sees several considerations as counting in favor of the same action, one cannot choose to act on one of them rather than another unless one downgrades some of the reasons by changing one's mind about whether they really do count (or count sufficiently) in favor of the action in question. Changing one's mind in this way is a judgment – a decision – but not a *choice*.⁸

⁶ Ibid.

⁷ Ibid., 58

⁸ Ibid., 60

Nevertheless, it seems intuitive that the particular reasons for which an agent performs a given action *can* make her blameworthy, as in *Rescue to Riches*. Thus, Scanlon maintains that blameworthiness and impermissibility “vary independently.”⁹

Let’s see if we can reconstruct this argument more precisely. Clearly Scanlon thinks that something about the relationship between permissibility and choice places constraints on what can be impermissible. I agree. Nevertheless, it isn’t quite right to say that, because questions of permissibility “typically” arise in deliberative contexts, that they apply only to “alternatives between which competent agents can choose.”

For one thing, the inference from an observation about when questions of permissibility usually arise to a *universal* limitation on their application is unlicensed. While it is true that we often consider questions of permissibility when we are trying to decide what to do, we also raise such questions post deliberation, regarding what either is being or has already been done. I may wonder, even while doing so, whether I am right to be editing my child’s essay for English class, perhaps affecting her final grade. And I may retrospectively doubt this decision too. To say that questions of permissibility apply only to “alternatives” that agents can choose makes it sound as if we can only apply the terms “permissible,” “impermissible,” and “obligatory” to *possible* courses of action arrayed before a deliberating agent. But of course those terms can qualify *actual* actions, which agents *are* choosing or have *already* chosen, too.

⁹ Ibid., 55

Additionally, I don't think the terms "impermissible" and its variants apply only to events that we bring about by choosing to, or to *actions*. *Failures* to bring events about by choosing to, or *omissions*, can be impermissible too. Yet, it isn't clear to me that agents must make *any* choice in order to impermissibly omit. I may break a promise to you, not by choosing to break it but simply by failing to keep it. Thus, for these two somewhat technical reasons I am dissatisfied with the claim that questions of permissibility apply only to alternatives between which agents can choose.

Nevertheless, I think Scanlon is quite right that the scope of the impermissible is in some way constrained by our capacity for choice-making. Questions of permissibility feel native to the deliberative context, because to call something "impermissible," *is* to call it a failure of choice. This is just what the term "impermissible," as opposed to, say, "bad" or "unfortunate," *means*. There are certain events to which questions of permissibility do not apply, because agents can neither bring them about by choosing to nor – as in cases of omissive *failures* to do as required – *prevent* them by choosing *not* to omit. Being born in Albuquerque is just not the sort of event that can be "impermissible" (or, for that matter, "permissible" or "obligatory") for an agent to bring about, because an agent's birthplace does not depend on her choices. Thus, I think the best way to render our constraint on the impermissible is this: An agent behaves impermissibly in the occurrence of some event, only when that event can either be brought about *or prevented* via her choices.

Having worked out the constraint that choice places on impermissibility, we are now in a position to reconstruct Scanlon's argument more precisely: (P1) An agent does something impermissible in the occurrence of some event only when that event is one that she is

capable of either bringing about or preventing via choice. (P2) Performing a given action for one reason as opposed to another is *not* the sort of event that agents are capable of bringing about or preventing via choice, because agents cannot choose the particular reasons for which they perform a given action. Agents cannot, for instance, choose to save a drowning person in order to aid someone in need as opposed to, say, choosing to save him in order to prevent his heir from inheriting. (C1) Therefore, an agent does nothing impermissible in performing a given action for one reason as opposed to another. (P3) But agents can be blameworthy for performing a given action for one reason as opposed to another. For instance, I can be blameworthy in saving a drowning person in order to prevent his heir from inheriting, but not in saving him simply to aid him. (C2) Therefore, agents can be blameworthy without behaving impermissibly.

I've just argued for P1, and the argument relies on the intuitive plausibility of P3. One *may* have a different intuition about this case, but I find it pretty plausible to hold that I am blameworthy. Thus, I don't want my argument to depend on disputing this intuition, and so I'll just accept P3. Plainly the remaining discussion to be had concerns P2. To be clear, Scanlon's view is not that agents *never* have control over the reasons for which they act. If an agent is facing a choice about which of two actions to perform, each of which is supported by different reasons, then Scanlon maintains that the agent can bring it about that he acts for the reasons that support the first action and not the second, by choosing to perform the first action and not the second.

That said, Scanlon does not think that agents can *choose* which considerations they see as reasons. Thus, he does not think that agents can, when faced with the prospect of

performing a particular action, control whether they perform *it* for one reason as opposed to another. After all, if we can't choose which considerations we see as reasons, then how could we bring about or prevent our performance of a particular action for one reason as opposed to another via choice? If I can't choose what I see as reasons, then I certainly can't choose to (say) save the drowning person in order to aid someone in need as opposed to saving him in order prevent his heir from inheriting by *choosing to see aiding him as my reason* for saving him. Nor could I *prevent* myself from saving him in order to stop his heir from inheriting, as opposed to (say) in order aid someone, by making a choice. For, to prevent myself from saving him in order to stop his heir from inheriting (while still, in fact, saving him), I would presumably have to choose to see something else – such as aiding him – as my reason. But, if I can't choose what to see as reasons, then that's just what I cannot do.

But perhaps we needn't agree with Scanlon that agents cannot choose what they see as reasons. Scanlon, we have seen, thinks they cannot, because he believes that seeing a certain consideration as a reason is less like choosing, and more like deciding what to believe. I am uncertain about this analogy. Nevertheless, I don't think my ambivalence warrants rejecting P2 because there is additional reason to doubt that agents can choose the reasons for which they act. Consider this variation on a doctrine of double effect case:

The Commander

Suppose you are prime minister, and the commander of the air force describes to you a planned raid, that would be expected to destroy a munitions plant and also kill a certain number of civilians, thereby probably undermining public support for the war. The commander asks you whether you think the raid is morally permissible and you – being a proponent of the doctrine of double effect – reply, “Well, that depends on what your intentions would be in carrying it out. If you intend to kill the civilians then the raid is impermissible, but if their deaths would be merely an unintended foreseeable side effect, it is permissible.” So the commander

*replies, "Alright, in that case I'll conduct the raid just in order to destroy the munitions plant," and goes ahead and bombs the plant.*¹⁰

If this is what happens, *has* the commander bombed the plant in order to destroy enemy weapons? It seems that, in attempting to select destroying enemy weapons as his reason for bombing the plant, the general has unwittingly not acted for that reason at all, but rather for a different reason altogether: he bombs the plant because he thinks doing so with certain intentions is permissible. Both your advice to the commander and his response are very strange here. I suggest that the strangeness comes from thinking of reasons for acting as *what* we choose when we make choices.

The problem is that the choice to perform some action (A) for some reason (R) would *itself* need to be made for a reason. But the reasons to A-for-the-reason-R are generally *not* R, because they are not generally considerations that support A-ing. Rather they are considerations that support *A-ing-for-the-reason-R*. And so long as R and the reasons to A-for-the-reason-R are not the same, then an agent's choice to A-for-the-reason-R will be ineffective, because in trying to make it, he won't be A-ing for the reason for which he is choosing to A (R), but rather for some *other* reason: namely, some consideration that supports *A-ing-for-the-reason-R*. That choosing to act-for-particular-reasons generally requires acting for *different* reasons than those that the agent is choosing, suggests that, in general, agents cannot choose to act-for-particular-reasons.

At this point I think we've raised a question not just for those, like me, who want to defend the impermissibility thesis, but for anyone who thinks an agent's motivating reasons *are*

¹⁰ This case is a variation on one of Scanlon's: (2008), 19-20

relevant to the permissibility of what she does. For Scanlon's argument now casts doubt on how it's *possible* for motivating reasons to be relevant to permissibility. What makes the difference between an agent's acting permissibly and her acting impermissibly seems to be what she *chooses* to do. Agents only behave impermissibly by choosing to do something they ought not to or failing to choose to do what they should. But if agents cannot choose the reasons for which they act, then how can an agent's motivating reasons turn an otherwise permissible action into an impermissible one?

IV. The Response

As perhaps you already know, Scanlon doesn't think that an agent's motivating reasons are *never* relevant to permissibility. To see where the argument goes wrong, it is helpful to examine a case in which he *does* take intent to be relevant to permissibility:

Rat Man

A man goes to the store to purchase some rat poison because he plans to put the poison in his wife's food and, thereby, kill her.

Is Rat Man's action of buying the poison itself impermissible? According to Scanlon, Rat Man's venomous purchase *is* wrong because, "in buying the poison the man is facilitating a murder."¹¹ His motivating reasons *do* affect the permissibility of his action "by making it part of a larger course of action that is impermissible."¹² Scanlon writes:

One thing we can say about a person who is buying poison with the intention of using it to kill his wife is that what he intends to do is impermissible, and that he should abandon that intention.¹³

¹¹ Ibid. 42

¹² Scanlon (2008), 41

¹³ Ibid. 41

Here, Scanlon suggests that Rat Man *can* control whether he acts for certain reasons, by abandoning his intention to kill his wife. If he chooses not to kill his wife, then his plan to do so will no longer *be* a reason for him to buy the poison. Of course, if Rat Man abandons his plan to kill his wife, then presumably he won't purchase the poison at all, since those plans were his only reason to, to begin with. Thus Scanlon is likely thinking that Rat Man will here control whether he acts for certain reasons by refraining from acting altogether. But why don't we alter the case's details so that our man sees multiple considerations counting in favor of buying the poison?

Rat Man II

Although the rat poison at issue is fatal to both rodents and humans, it is actually medically useful for dogs. Furthermore, Rat Man's dog, Robin, is suffering from a painful and, potentially fatal, stomach ulcer, which, his vet has informed him, only this particular brand of rat poison can cure.

From Rat Man's point of view there are now two sufficient reasons to buy the poison. He can use it to kill his wife or to cure his dog. Now the choice with which our man is faced is whether to perform a particular action for one reason as opposed to another. We may suggest that what he *should* do is buy the poison in order to save Robin but not in order to kill his wife. But we *can't* say this if we are to treat this case as Scanlon must. For now Rat Man is faced with the choice to perform a particular action supported by multiple reasons, and so P2 kicks in: agents can't choose the particular reasons for which they perform given actions. Additionally, if *why* Rat Man purchases the poison isn't up to him, then it can't be impermissible to buy it to kill his wife, but permissible to buy it to save his dog. And yet, Rat Man still *ought* to purchase the poison, even if he can't choose *why* to. After all, he had better save Robin. Thus, it must be permissible for him to purchase the poison regardless of his reasons for doing so! He may, of course, be *blameworthy* if he makes the purchase because of

his plans to kill his wife. But now the case just looks like another counterexample to the impermissibility thesis. Rat Man will be blameworthy, while doing nothing impermissible.

The problem with treating the case this way is that bombarding Rat Man with additional reasons to purchase the poison just shouldn't affect whether he can abandon his uxoricidal plans or not. But if he can do this in the case's second incarnation, then he *can* bring it about via choice that he purchases the poison for one reason as opposed to another. If he chooses not to kill his wife, then, just as in *Rat Man I*, his plan to kill his wife will no longer *be* a reason for him to purchase the poison. But since, in *Rat Man II*, purchasing the poison will still help his dog, our man *will* still purchase the poison in order to help his dog. Thus, he can bring it about that he purchases the poison not in order to kill his wife, but rather only in order to cure his dog, *by choosing not to kill his wife*.

Scanlon's argument against the impermissibility thesis is therefore unsound because its second premise turns out to be false. Agents *can* bring it about via choice that they perform particular actions for some reasons as opposed to others. But P2 certainly *seems* compelling, which is why the argument presents an interesting challenge to the impermissibility thesis. P2 seems compelling because, although agents *do* bring it about via choice that they perform actions for the particular reasons that they do, they don't do so *by* thinking about the various reasons for which they might perform a given action and then choosing to be motivated by some of those reasons rather than others. Rather, they bring it about that they perform the actions that they do for the particular reasons that they do simply *by choosing what to do*. In purchasing the poison in order to kill his wife, Rat Man chooses to kill his wife. Thus, he certainly acts impermissibly in purchasing rat poison for that reason, because purchasing the

poison for the reason that he does *is* something his brings about via choice! But he does so by choosing *to kill his wife* (and to purchase rat poison as a means of doing so) – not by choosing to see killing his wife as his reason for purchasing the poison.

Incidentally, I say that Rat Man chooses to kill his wife “in” acting for the reason that he does, as opposed to saying that so-choosing is “necessary for” acting for the reason that he does, because I think the connection between his acting for the reason that he does and his choosing to kill his wife is very tight. It’s not simply that his acting for the reason that he does indicates that he has *already* chosen to kill his wife. Rather, his acting for the reason that he does *partially constitutes* his choice to kill his wife. To be motivated by the aim of killing one’s wife, my thought is, *is what it is* to choose to kill one’s wife, and in allowing himself to be so motivated while purchasing the poison, he continues to make that choice.

Can agents choose what they see as reasons then? It depends what we mean. If we mean, “Can an agent first decide to perform an action, and then decide to regard certain considerations as her reasons for doing so?” then the answer is “No.” But if we simply mean, “Can agents control which considerations they see as reasons through their capacity for making choices?” then the answer is “Yes.” It’s just that we control which considerations we see as reasons not by attempting to will ourselves to be motivated by particular considerations, but rather simply by choosing which aims to pursue. And because we control which considerations we see as reasons via our choices, we do bring it about that we perform given actions for certain reasons and not others via our choices too.

Let's accept that P2 is false, and that agents can bring it about via choice that they perform particular actions for some reasons and not for others. What of the initial "right thing for the wrong reason" case, *Rescue to Riches*? Even if I *can* bring it about that I act for some reasons as opposed to others by making choices about what to do, what can I choose to do in this case so that I save the drowning person, but *not* merely to prevent his heir from inheriting? What I can do, as others have suggested, is choose not merely *to* save the drowning man, but also the *conditions* under which I am prepared to save him.¹⁴

If, as stipulated, I would be happy to see the drowning man die, and am able to bring myself to save him *only* in order to prevent his heir from inheriting, then I appear to be committed to following through with the rescue *only* provided certain conditions obtain. As described, I seem committed to *ceasing* my rescue attempt if I find out our drowning victim has lost all his money in the stock market, or has written his presumptive heir out of his will. But if this is right, then the *choice* I make in pursuing my course of action is to rescue the drowning man *only provided that doing so is a necessary means of preventing his heir from inheriting*. But this is not the choice that I ought to make. (Rather, I should choose to rescue him *regardless of whether doing so is a necessary means to preventing his heir from inheriting*.) Thus, in acting for the reasons that I do, I *am* making a choice I ought not to, and, thereby, acting impermissibly: I am pursuing a plan to rescue a drowning victim only provided that doing so is a necessary means to preventing his heir from inheriting. Consequently, although my motivating reasons do render me blameworthy, this case is no counterexample to the impermissibility thesis.¹⁵

¹⁴ Kolodny (2011), 105. I have tried to stress here a point that I think is left largely implicit in Kolodny's discussion: that forming an intention (conditional or otherwise) and, thereby, controlling one's reasons for acting requires not *choosing to act for a particular reason*, but rather, *choosing what to do*.

¹⁵ One could interpret the case slightly differently. As I've interpreted it here, I am committed to *not* rescuing the drowning man if doing so ceases to be necessary for preventing his heir from inheriting. But what if,

But surely saving a drowning person is obligatory – not impermissible! If I claim that I act impermissibly in this case, then aren't I also unacceptably claiming that saving a drowning person is impermissible?

No. There are two ways of describing this case without falling prey to this result. Pick your favorite: First, one might reply that although saving a drowning person is, of course, permissible, “saving a drowning person” is not an entirely accurate description of what I am doing in the present case. Rather, the *accurate* description of what I am doing is “saving a drowning person only under the condition that doing so is a necessary means to preventing his heir from inheriting,” which is impermissible. Alternatively, one might reply that there are many descriptions of what I am doing in the present case, under some of which I am acting permissibly and some of which I am not. “Saving a drowning person” is a description of what I am doing under which I am acting permissibly, and, indeed, it *is* permissible for me to save a drowning person, just as I am doing. However, “pursuing a plan to save a drowning person only under the condition that doing so is a necessary means to preventing his heir from inheriting” is another accurate description of what I am presently doing, and it is impermissible for me to do this. Thus, although saving a drowning person is permissible, it

instead, we read the case as one in which I am committed to saving him insofar as doing so *is* a means of preventing his heir from inheriting, but just *haven't committed* one way or the other regarding what I will do if rescuing him ceases to be necessary to this end? On this alternative interpretation, my intuitions regarding impermissibility begin to fail me. However, this does not show that blameworthiness and impermissibility come apart. Rather, I think that *whether* I am blameworthy depends on *whether* I act impermissibly in choosing to rescue a man in order to prevent his heir from inheriting without *also* making additional commitments regarding other conditions under which I will save him – a question on which, I, personally, am ambivalent. I mentioned earlier that one might have different intuitions regarding my blameworthiness in this case. Part of the reason why this is so, I think, is precisely that, on some interpretations of it, whether or not what I do is impermissible becomes less clear.

does not follow that I am doing nothing impermissible, because another thing I happen to be doing is carrying out a nefarious plan.

I claim that, in *Rescue to Riches*, the *conditions* of performance specified by my choice to act are what make my action wrong. This may make you worry that calling this a case of impermissible behavior is just a “bit of a stretch.” Perhaps, moreover, you think this weighs against my treatment of it. I have two replies to this worry.

First, it is certainly no *oddity* that a choice to act may involve undertaking commitments regarding the conditions of one’s performance. For instance, one may set out with his lanterns to warn of the Redcoats, planning to hang “One if by land; two if by sea.”¹⁶ Nor is it particularly anomalous that the conditions under which one commits to following through with an action might make a difference to that action’s permissibility. When Trump chooses to put “America First” in foreign policy, part of what is wrong this choice surely consists in its commitment to *never* consider the needs of non-Americans over Americans, no matter how dire non-American needs might be.

My second response to the “bit of a stretch” objection is to remind you of my current dialectical position. I already made my positive case for the impermissibility thesis in section II of this paper. There, I argued that we ought to believe the impermissibility thesis because it is the principled explanation of our blameworthiness intuitions. If *Rescue to Riches* were a counterexample to the impermissibility thesis, this would of course give us reason to reject that thesis in spite of my positive argument. But given that I have offered an independent

¹⁶ I take this example from Gibbard (2003), 54.

positive argument for the impermissibility thesis, I take my burden in responding to putative counterexamples to be to show simply that any proposals can, in fact, be reinterpreted in a manner that reveals them *not* to be counterexamples. My burden is *not* to show that putative counterexamples can be explained away *effortlessly*.

My response to *Rescue to Riches* and to *Rat Man II* do both employ what we can now regard as a general strategy for replying to any putative “right thing for the wrong reason” cases proposed as counterexamples to the impermissibility thesis. An agent’s reasons for acting generally indicate something about what she is committing herself to bringing about in acting: what she is choosing to do under what conditions, or what *plans* she choosing to pursue. Thus, in any case where an agent is acting for dubious reasons, those reasons will likely indicate that she is either pursuing plans she ought not to be, or failing to pursue plans that she should. In making or failing to make such a choice, she will be behaving impermissibly. So, for any case in which an agent appears blameworthy for acting for bad reasons, I think we can show that she is behaving impermissibly in doing so.

V. Conclusion

In this paper, I first argued that we ought to accept the impermissibility thesis because it is the principled explanation of our blameworthiness intuitions. Then, I rebutted Scanlon’s argument against the impermissibility thesis by showing that putative “right thing for the wrong reason” cases are not actually counterexamples to it.

But perhaps you are wondering whether “right thing for the wrong reason cases” were really the most threatening cases for the impermissibility thesis to begin with. Maybe the most

difficult cases are those in which agents are potentially blameworthy not for *acting*, or even for acting for particular *reasons*, but rather *simply* for their attitudes. One might think, for example, that individuals can be blameworthy just for failing to *care enough* about others, without such attitudes ever manifesting themselves in impermissible behavior.

I lack the space to treat such cases extensively here, but let me at least gesture at my response: Many cases in which agents are blameworthy for their attitudes are cases of indirect blameworthiness. Individuals can be blameworthy for failing to care enough, for instance, when such failures are the reasonably foreseeable consequences of their earlier impermissible behavior. But I have also now demonstrated, in responding to Scanlon, that taking a managerial stance towards our attitudes is not the only way we have of controlling them. Another way in which we control at least some of our attitudes – such as the reasons for which we act – is by *constituting* them via our choices about what to do. And so I think that agents can also be *directly* blameworthy for their attitudes, when and because those attitudes are constituted by their making or failing to make choices. Sometimes all it *is* to care about something or someone is to choose to conduct oneself with respect towards that thing or person, and all it is to fail to care is to omit to so-choose. But I discuss these more difficult cases in further detail elsewhere. Here, I've just made an initial case for the impermissibility thesis, and defended that thesis against an objection.

References

- Capes, J. (2012). Blameworthiness without Wrongdoing. *Pacific Philosophical Quarterly*, 93(3), 417–437.
- Fischer, J. M., & Ravizza, M. (1998). *Responsibility and Control: A Theory of Moral Responsibility*. Cambridge University Press.
- Fischer, J. M., & Tognazzini, N. (2012). The Triumph of Tracing. In *Deep Control: Essays on Free Will and Value*. Oxford: Oxford University Press.
- Gibbard, A. (2003). *Thinking How to Live*. Cambridge, MA: Harvard University Press.
- Kolodny, N. (2011). Scanlon's Investigation: The Relevance of Intent to Permissibility. *Analytic Philosophy*, 52(2), 100–123.
- Scanlon, T. M. (2008). *Moral Dimensions: Permissibility, Meaning, Blame*. Cambridge, MA: Harvard University Press.
- Scanlon, T. M. (2013). Interpreting Blame. In D. J. Coates & N. Tognazzini (Eds.), *Blame: Its Nature and Norms* (pp. 84–99). Oxford University Press.
- Smith, A. (2013). Moral Blame and Moral Protest. In D. J. Coates & N. Tognazzini (Eds.), *Blame: Its Nature and Norms* (pp. 27–48). Oxford University Press.
- Wolf, S. (2011). Blame, Italian Style. In R. J. Wallace, R. Kumar, & S. Freeman (Eds.) (pp. 332–347). Oxford University Press.