Dilemmas, Luck, and the Two Faces of Morality

Johann Frick (jfrick@berkeley.edu)

1. Introduction

According to some philosophers, there are *moral dilemmas*: situations where, no matter how an agent chooses to act, her action will be morally wrong. According to some philosophers, there is *moral outcome luck*: how blameworthy an agent was for performing an action can depend on features of the action – namely how it turned out – that were at least partly beyond the agent's control. Over the past few decades, much ink has been spilt on both topics, often by the same authors. Yet the two issues are rarely connected. I want to argue that this is an oversight. We can gain a deeper understanding of both problems by examining them in tandem.

The argument of this paper unfolds two stages. Sections 2 and 3 motivate my parallel investigation by examining the features that make a moral theory admit of moral outcome luck, on the one hand, and of moral dilemmas, on the other. My conclusion is that there is a connection between these seemingly disparate normative phenomena. Proponents of moral dilemmas and proponents of moral outcome luck both implicitly endorse a thesis I call

> **Parochial evaluation:** In evaluating a given action a_i or the agent who performed it, it is sometimes appropriate to "screen off" portions of modal space or at least to "downplay" their evaluative significance. In the case of moral outcome luck, this is the space of *alternative possible outcomes* $\{o_1; o_2; \ldots o_n\}$ that the action could have resulted in. In the case of moral dilemmas, this is the space of *alternative possible actions* $\{a_1; a_2; \ldots a_n\}$ that were available to the agent. Instead, we should make our assessment of the action or the agent exclusively or at least primarily a function of facts about a_i or o_i themselves.[1]

---

[1] This is not, of course, to claim that any moral view that admits of moral dilemmas also admits of moral outcome luck, or vice versa. It is only to claim that moral theories that give rise to either phenomenon have a structural commonality, namely a commitment to parochial evaluation.

Proponents of moral outcome luck, as we shall see in Section 2, are parochialists about *outcomes*. They hold that an agent's degree of blameworthiness for performing a risky action a_i is not primarily a function of what possible outcomes $\{o_1; o_2; \ldots o_n\}$ the action *could* have resulted in. Rather, what matters crucially are facts about the outcome o_i which *actually* results.[2]

Proponents of moral dilemmas, as we shall see in Section 3, typically subscribe to parochialism about *options*. Moral dilemmas arise when we decouple the question of whether a given course of action a_1 is wrong from the wider context of choice – what alternative courses of action the agent could have chosen instead in that situation. Rather, a_1 is wrong just because of what it is *in itself*. That the only alternative to a_1 is a_2 – which is *also* terrible – does not make it any less the case that it is wrong to do a_1. Apply the same reasoning to a_2 and we arrive at a moral dilemma.

In the remainder of the paper, I then turn to two substantive ethical questions: should we accept that there are moral dilemmas, and should we accept that there is moral outcome luck? There are well-known and powerful objections to answering either question in the affirmative. Embracing moral outcome luck seems to fly in the face of the powerfully intuitive Control Principle, according to which an agent cannot be morally assessable for features of his action that were not under his control.[3] And accepting that there are genuine moral dilemmas has struck many philosophers as extremely unpalatable, both on formal and on moral grounds.

Despite these misgivings, some philosophers have felt compelled to concede the possibility of genuine moral dilemmas and of moral outcome luck, for reasons that are, again, instructively parallel. Only by admitting moral dilemmas and moral outcome luck, they argue, can we make sense of certain aspects of our ordinary ethical experience – for instance, the very different *reactive attitudes* that seem appropriate in purported cases of moral outcome luck, depending on how things turn out[4], or the sense that in cases of tragic moral conflict, no matter how the agent acts, it will be appropriate to feel guilt

---

[2] Other forms of moral luck, such as "circumstance luck" or "constitution luck" are not discussed in this paper.

[3] Thomas Nagel, "Moral Luck" in his *Mortal Questions* (New York: Cambridge University Press, 1979).

[4] Bernard Williams, "Moral Luck" in *Proceedings of the Aristotelian Society*, Supplementary Volumes, Vol. 50 (1976), pp. 115-135 and Nagel, "Moral Luck" (op. cit.).

for what he has done. [5] I refer to these, collectively, as *arguments from phenomenology*, since they maintain that only by allowing the possibility of genuine moral dilemmas and genuine outcome luck can we ratify the characteristic moral emotions that seem apt in such cases.

Section 4 introduces a novel strategy for resisting such arguments. The key that unlocks the twin problems of moral dilemmas and moral outcome luck, I argue, lies in clearly differentiating an *agent-* and a *patient*-centric perspective of evaluation – what I call the "two faces of morality". From the agent-centric perspective, the central question is "Has the agent comported herself well or poorly, morally speaking?" Facts about the rightness or wrongness of the agent's action belong to this first, agent-centric perspective. Also, on some views of blameworthiness, questions about the agent's degree of blameworthiness are exclusively the province of this agent-centric perspective of evaluation. [6] By contrast, the central question from the patient-centric perspective is "How have people been *affected* by the agent's action? Specifically, has anyone been *wronged* by the agent's behavior? If yes, how seriously?"

I then argue that, bracketing for the moment the question whether we should accept parochial evaluation with regard to *agent-centric* questions about moral rightness or blameworthiness, it is extremely plausible that when it comes to the *patient*-centric question whether someone has been *wronged* by the agent's action (and if so, how seriously), we *should* be parochialists, both (i) about outcomes and (ii) about options. As I show in Section 5, how seriously a person has been wronged by an action depends to a large degree, not on what outcome the action *could* have produced, but on what outcome *actually* occurred. And as I argue in Section 7, on the back of an account of wronging which I develop in Section 6, there are tokens of certain extreme act-types (killing, torture, rape) which can be said to wrong a patient, *no matter* what alternatives the agent had to treating her in this way.

Once we realize that we must be parochialists about *wronging*, however, we can use this insight to solve the problem of moral dilemmas and defang the

---

[5] See, for instance, Bernard Williams, "Consistency and Realism," *Proceedings of the Aristotelian Society* (Supplement), 40 (1966), pp. 1–22; Michael Walzer, "Political Action: The Problem of Dirty Hands," *Philosophy and Public Affairs* Volume 2 (1972), pp. 160–180; Ruth Barcan Marcus, "More about Moral Dilemmas," in H.E. Mason (ed.) *Moral Dilemmas and Moral Theory* (New York: Oxford University Press, 1996), pp. 23–35.

[6] More on this question in Section 9.

problem of moral luck. I argue, in Section 8, that the distinctive phenomenology of tragic moral conflicts can be fully accounted for in terms of parochialism about *wronging*. There is no need to postulate genuine moral dilemmas in order to explain why, in cases of tragic moral conflict, an agent can rightly believe that, no matter how he acts, he will have reason to feel guilty. With regard to the problem of moral outcome luck, I show in Section 9 that what accounts for the different reactive attitudes that seem appropriate vis-à-vis lucky and unlucky agents are not *agent-centric* facts about how well or poorly the agent has comported herself, but *patient-centric* facts about how seriously other people have been *wronged* by the agent. Depending on the account of blame we endorse, this implies that either there is no moral luck or, if there is, it is not of a kind that conflicts with the Control Principle.

Section 10 concludes with some broader reflections on the significance of the two faces of morality for our ethical life.

2. Moral outcome luck and parochialism about outcomes

Many of the canonical illustrations of moral outcome luck juxtapose two agents who each, without adequate justification, perform some risky action. Consider the following case, after Thomas Nagel:

> *Drunk Drivers:* Lara and Ursula are two drunk drivers who go on the road, equally intoxicated, with the same blithe indifference to the safety of other motorists and passers-by, and exposing others to equal levels of risk. However, Ursula is unlucky and runs over a child, whereas lucky Lara 'gets away with it' and returns home without incident.

By assumption, Ursula and Lara's actions had the same profile of *possible* outcomes, with identical likelihoods attached to each outcome. To keep things simple, let us suppose that there were only two relevant ways their actions could turn out: {Get home without incident (prob. 0.98); have an accident and kill a person (prob. 0.02)}. The crucial difference, of course, is where in this modal space Lara and Ursula actually end up.

To endorse moral outcome luck is to endorse the position that Lara and Ursula's levels of blameworthiness are *not* just a function of features of their

behavior that were, in some sense, up to them: their mental states and the risk that they imposed on other people. Rather, to a significant degree, the agents' blameworthiness is a function of something they don't control, namely how their risky actions turn out. Ursula is more blameworthy than Lara, it is claimed, because she actually *killed* a person. "True enough," says the proponent of moral luck, "Ursula's drunk driving *could* have turned out just as benignly as Lara's. But fatefully, it didn't. That makes all the difference."

Proponents of outcome luck thus embrace a form of parochial evaluation – what we might call "parochialism about *outcomes".* Compared to an opponent of moral luck, a proponent of outcome luck, in determining Lara and Ursula's respective levels of blameworthiness, downplays the significance of facts about *possible* outcomes (that Lara and Ursula's actions had the same likelihood of turning out well, or of ending in disaster) and gives disproportionate weight to their action's *actual* outcomes.

This is not to say that friends of moral outcome luck need give *no weight at all* to the alternative outcomes that an agent's action could have had. Parochial evaluation, taken to such an extreme, would surely be implausible. Consider reckless Rita, who goes on the road, not just mildly intoxicated like Lara and Ursula, but *stinking drunk*, thereby putting people at a far greater risk of accident. (The space of possible outcomes for Rita, let us suppose, was {Get home without incident (prob. 0.60); have an accident and kill a person (prob. 0.40)}. However, like Lara, but unlike Ursula, Rita makes it home without causing injury to anyone.

Intuitively, Rita is more blameworthy than Lara. Neither agent ended up actually causing harm, but Rita put people at far greater risk. That modal fact surely matters to her level of culpability.

Crucially, however, the proponent of moral luck can maintain that although Rita is more blameworthy than Lara, she is *less* blameworthy than unlucky Ursula. Although Rita put people at far greater risk than Ursula, the fact that Ursula's behavior actually ended in disaster, whereas Rita had a lucky escape, has an outsize effect on their respective levels of blameworthiness. If we deny the possibility of moral outcome luck, by contrast, we are committed to viewing Ursula as less blameworthy than Rita. (For, in that case, Ursula is no more blameworthy than Lara; and Lara, we said, is less blameworthy than Rita).

Let us turn next to moral dilemmas.

3. Moral dilemmas and parochialism about options

Following Thomas Nagel, I define a moral dilemma as a situation "which leaves [an agent] with nothing to do which is not wrong".[13] We can distinguish two varieties of moral dilemmas: in *prohibition dilemmas* each feasible course of action[14] is prohibited. In *obligation dilemmas*, more than one feasible course of action is obligatory, such that failing to take it would be morally wrong, yet the obligatory courses of action are not *compossible*, i.e. cannot *all* be feasibly performed in the situation.[15] There can be mixed cases as well.

Consider two stock examples from the literature, which some philosophers have characterized as moral dilemmas:

> *Ticking Time Bomb:* A terrorist has hidden a nuclear device in downtown Boston. The device is on a timer, and if the bomb is not defused within the next 10 minutes, it will detonate, killing 500,000 people. The agent knows that the only way to get the terrorist to reveal the location of the bomb is for him to torture the terrorist's five year old daughter in front of him. (Imagine that the terrorist and his daughter are being detained in New York City, and wouldn't themselves perish in the explosion).[16]

---

[13] Thomas Nagel, "Moral Luck" in his *Mortal Questions*, p. 70, fn. 9. See also Bernard Williams, who understands moral dilemmas, or 'tragic moral conflicts', as situations where "an agent can justifiably [and correctly] think that whatever he does will be wrong." Williams, "Conflicts of Values" in *Moral Luck,* (Cambridge: Cambridge University Press, 1981), p. 74. Nagel's definition of moral dilemma is stronger than that adopted by some other authors, for instance Walter Sinnott-Armstrong, who understands moral dilemmas as "any situation in which the agent cannot avoid violating a non-overridden moral requirement". Sinnott-Armstrong's definition is less demanding than Nagel's since, as he clarifies, one can violate a non-overridden moral requirement without *eo ipso* acting impermissibly. See Walter Sinnott-Armstrong, "Moral Dilemmas and Rights" in H.E. Mason (ed.), *Moral Dilemmas and Moral Theory* (New York: Oxford University Press, 1996), pp. 48–65.

[14] A "course of action", as I employ the term, may correspond to an action or an omission.

[15] For further discussion of the distinction between prohibition and obligation dilemmas, see Peter Vallentyne, "Prohibition Dilemmas and Deontic Logic," *Logique et Analyse*, Vol. 30 (1987), pp. 113–122 and "Two Types of Moral Dilemmas," *Erkenntnis*, Vol. 30 (1989), pp. 301–318.

[16] Adapted from a similar case in Michael Walzer's, "Political Action: The Problem of Dirty Hands", *Philosophy & Public Affairs*, Vol. 2, No. 2 (Winter, 1973), pp. 160-180. In Walzer's version of the case, it is the terrorist himself who must be tortured to reveal the location of the bomb. My modification to Walzer's case makes it, if anything, a stronger candidate for a moral dilemma. If, as Walzer believes, it would be morally wrong to torture the terrorist – despite the fact that he is

In this supposed prohibition dilemma, the agent must choose between two courses of action:

> A: Torture the terrorist's daughter.
> B: Allow 500,000 people to perish in a nuclear explosion.

The second example is

> ***Sartre's Student:*** A young student, whose brother had been killed in the German invasion of France in 1940, wants to join the Free French forces, to avenge his brother and fight the Nazi occupation. Yet doing so would mean leaving behind his ailing mother, who is living with him and is heavily dependent on him, both physically and for emotional support.[17]

This is a supposed obligation dilemma. Since the obligatory courses of action are not compossible, the student cannot but violate one of his putative obligations. He must either

> Y: Abandon his mother.
> Z: Give up the fight against the Nazi occupation.

Philosophers who view these cases as genuine dilemmas (e.g. Thomas Nagel and Michael Walzer for the first case[18] and Bernard Williams and Ruth Barcan Marcus for the second[19]) hold that, in each vignette, *both* courses of action are

---

the source of the unjust threat and could avoid torture by revealing the location of the bomb (which he has independent moral reason to do) – then *a fortiori*, it must be wrong to torture the terrorist's innocent daughter, who has no responsibility for creating the threat and who could not escape torture by revealing the location of the device. A structurally similar case in which an innocent third party must be tortured in order to prevent a great evil is discussed by Thomas Nagel, *The View from Nowhere* (New York: Oxford University Press, 1986), Chapter IX.

[17] After Jean-Paul Sartre, "Existentialism is a Humanism," trans. Philip Mairet, in Walter Kaufmann (ed.), *Existentialism from Dostoevsky to Sartre* (New York: Meridian, 1957), pp. 287–311.

[18] Thomas Nagel, "War and Massacre", reprinted in his *Mortal Questions* (Cambridge: Cambridge University Press, 1979), pp. 53–74 and Michael Walzer, "Political Action: The Problem of Dirty Hands" (op. cit.).

[19] Ruth Barcan Marcus,"Moral Dilemmas and Consistency", *The Journal of Philosophy*, Vol. 77 (1980), pp. 121–136.

morally wrong. This is not to say that they necessarily believe that the agent's practical quandary is unresolvable. Walzer, for one, suggests that the "correct" decision in the first scenario may be to torture, and that we should want our political decision-makers to be willing to "get their hands dirty" by choosing accordingly. However, this does not alter the fact that *whichever* alternative the agent chooses in this situation, he chooses a course of action that is morally wrong, all things considered.

I submit that moral theories that admit the possibility of moral dilemmas do so because they subscribe to a species of parochial evaluation – what I will call "parochialism about *options*". To see this, consider first how moral theories *avoid* giving rise to moral dilemmas.

A moral theory which rules out the possibility of moral dilemmas holds that in any choice situation, defined by a set of feasible alternatives, there is at least one permissible course of action for an agent to perform – however unattractive the options may be, individually considered. Call this the **"There is Always a Permissible Action"** principle.

Accepting this principle commits us to a certain way of determining whether a given action is wrong: Consider an agent who must decide between two available courses of action, a_1 and a_2. Would it be morally wrong for the agent to perform a_1? If we accept the above principle, it can't be that both available courses of action are morally wrong. In determining the deontic status of a_1, the agent must therefore consider the alternative, a_2, to see whether it has a stronger claim to being considered morally wrong. If it does, then a_1 cannot be morally wrong, on pain of violating the There Is Always a Permissible Action principle.

Rejecting the possibility of moral dilemmas thus commits us to what I shall call

> **Comparativism about moral wrongness:** Whether a given course of action can be deemed morally wrong is never just a matter of its non-comparative properties; rather, it is partly a function of, and constrained by, what the feasible alternatives are.

Although this principle is not often made explicit, it is implicitly respected by various ways of analyzing situations like *Ticking Time Bomb* and *Sartre's Student* on which they do not come out as moral dilemmas.

Suppose the agent in *Ticking Time Bomb* chooses option A, and tortures the terrorist's child to save Boston from obliteration. Was this morally permissible? Here are two possible answers to this question:

Proponents of *lesser evil reasoning* answer 'Yes'. Given that the agent's only alternative to A was B, it was permissible to do A. A was the "lesser evil". Lesser evil reasoning is, of course, a corollary of act-consequentialism; but one need not be a consequentialist to embrace it. *Threshold deontologists* differ from act-consequentialists by, among other things, positing that we are constrained in our pursuit of the good by deontological side constraints, for instance a constraint against torturing an innocent person. These constraints, however, are not *absolute*; they have thresholds beyond which they may permissibly be infringed, if doing so would produce a much greater good or prevent a much greater evil.

By contrast, proponents of what I will call *consistent absolutism* answer 'No'. They hold that the constraint against torturing an innocent person cannot, under any circumstances, be overridden by considerations of the *greater good* (at least for other people).[20] Given that doing A violated this constraint, the agent acted wrongly in doing A. It would have been right to choose B instead. Although the consequences of B are terrible, given that the only alternative to doing B was to do A, doing B would have been permissible. (Note that, lest it give rise to moral dilemmas, an absolutist position would have to maintain that, in a forced choice between violating a side-constraint vis-à-vis one person or violating a side-constraint vis-à-vis a different person, it is permissible to violate one or the other side constraint – perhaps the less stringent one. Insofar as absolutist moral systems do not contain this escape hatch, they will be prone to moral dilemmas for exactly the reason that I am about to describe. I will use "*consistent* absolutist" as a term of art for those variants of absolutism which avoid moral dilemmas).

---

[20] It is compatible with consistent absolutism, as I have characterized it, that although the moral side constraint against torturing an innocent person cannot permissibly be overridden just because doing so would bring about a much better outcome *for other people*, it might be permissible to torture a person in order to spare that *same* person from an even worse harm, or an even more serious violation. (Suppose that, compelled by a villain, I must either torture an innocent child, or torture *and then kill* that same child).

Despite reaching opposite verdicts about our case, lesser evil reasoning and consistent absolutism have something important in common. Both positions respect comparativism about wrongness.

The proponent of lesser evil reasoning readily admits that torturing an innocent child is a *terrible* thing to do. But given that the only alternative is, in Robert Nozick's phrase, the "catastrophic moral horror" of letting an entire city perish, he concludes that it is not, in the final analysis, wrong to resort to torture. Were the circumstances different and there existed, besides A and B, a third feasible course of action

> *C*: tricking the terrorist into divulging the location of the bomb,

then choosing A would, of course, be wrong. Thus, while the fact that A involves torturing an innocent child suffices to render it morally wrong, if the set of feasible options is {A; B; C}, this non-comparative property of A does *not* suffice to render it morally wrong if the choice set is restricted to {A; B}.

Likewise, the consistent absolutist concedes that allowing 500,000 people to perish is unspeakably awful. But given that the only available alternative is to torture an innocent child, it is not, in the final analysis, morally wrong. Again, were the set of feasible alternatives also to contain C, then choosing B would, of course, be wrong. Thus, while the fact that B represents a catastrophically bad outcome suffices to render it impermissible if the option set is {A; B; C}, it does not when the feasible courses of action are limited to {A; B}.

By contrast, the mark of moral theories that give rise to moral dilemmas is precisely that they do *not*, in this way, let ascriptions of moral wrongness be constrained by facts about the available alternatives. Typically, such theories admit moral dilemmas because they reject comparativism and instead embrace

> **Non-comparativism about moral wrongness:** There are some courses of action a_i which are morally wrong, *whatever* the alternatives available to the agent. The agent need not compare a_i to the available alternatives to determine that a_i is morally wrong; rather, the agent can know that a_i is wrong just in virtue of its non-comparative properties.[21]

---

[21] This is not to claim that embracing non-comparativism about moral wrongness is *logically* the only way in which a moral theory can generate dilemmas. I only claim that, when we examine the most convincing candidates for moral dilemmas in the literature and the types of moral

If we accept non-comparativism about wrongness, moral dilemmas can arise in at least two ways:

Firstly, there can be choice situations where *all* available courses of action have non-comparative properties that make them morally wrong, according to the non-comparativist – for instance because they violate an *absolute side-constraint* or fail to fulfill an *absolute duty*.

If the case of *Sartre's Student* does indeed constitute a bona fide moral dilemma, this is what is going on here. A non-dilemmatic analysis of this case might hold that what we observe is a clash of two Rossian *prima facie* obligations – say an obligation of "filial piety" to look after one's mother, and a "patriotic" obligation to liberate one's country from a vile occupier. But, on the Rossian analysis, when two prima facie obligations conflict, at most one of them can be the agent's *actual* obligation, which it would be morally wrong not to fulfil. The other prima facie obligation is no obligation at all.[22]

By contrast, those philosophers, such as Ruth Marcus or Bernard Williams, who regard cases like *Sartre's Student* as genuinely dilemmatic do so because they reject the Rossian analysis in terms of *prima facie* obligations, and view *both* obligations as *actual* obligations, which it is wrong not to fulfil. As Williams puts it: "The discovery that my factual beliefs conflict eo ipso tends to weaken one or more of the beliefs; not so (…) I think, with one's conflicting convictions about what one ought to do."[23]

The second way in which embracing non-comparativism can generate moral dilemmas is if some of the available courses of action are deemed to be wrong, simply in virtue of their non-comparative properties, whereas others are wrong because they violate some important comparative principle, such as a principle that tells us to always avoid the (much) greater of two evils.[24]

---

reasoning that underlie them, it is a commitment to non-comparativism about moral wrongness which seems to explain why we encounter a purported moral dilemma.

[22] W.D. Ross, *The Right and the Good* (Oxford: Oxford University Press, 1930), p. 30. Ross does concede that, faced with a conflict of prima facie duties, we often cannot know for certain how we ought to act. But the problem, here, is purely *epistemological*, not *ontological*. It does not reflect an underlying moral reality where there are two *actual* obligations that conflict.

[23] Bernard Williams, "Ethical Consistency", *Proceedings of the Aristotelian Society* (Supplement), 39 (1965), p. 109.

[24] The injunction to 'avoid the greater evil' is comparativist since, in order to know which, if any courses of action are wrong under this principle, we need to know the available alternatives. If

This is exactly the reasoning that leads Nagel and Walzer to characterize cases like *Ticking Time Bomb* as moral dilemmas. Nagel and Walzer strongly feel the pull of both types of moral consideration that animate the proponent of lesser evil reasoning and the moral absolutist – in Nagel's words, "a concern with what will *happen*" and "a concern with what one is *doing*", respectively.[25] Walzer, in particular, fully acknowledges the force of 'what will happen', going so far as to say that the "correct" decision (at least if the agent is a politician or responsible office holder) would be to prevent the greater evil, and that it would be morally wrong not to "get one's hands dirty".

Crucially, however, this does not cause Walzer to withdraw his claim that it is wrong to torture in this case. As Walzer writes, "one's hands get dirty from doing what it is wrong to do."[26] If the politician opts for torture, "it does not seem enough to say that he should feel very badly. … When he ordered the prisoner tortured, he committed a moral crime and he accepted a moral burden. Now he is a guilty man."[27] For Walzer, torturing a person is wrong, just because of what it is in itself, and regardless of the alternatives.

Moral theories which give rise to moral dilemmas thus typically do so because they accept non-comparativism about wrongness. And non-comparativism about wrongness is a species of *parochial evaluation*, namely parochialism about options. At least for certain courses of action, the non-comparativist believes that the modal context of available alternatives is wholly irrelevant. In evaluating such courses of action, we can screen off the space of alternative possible options, and focus solely on this course of action itself. As someone who believes that *Ticking Time Bomb* is a moral dilemma might say: "I don't *care* that your only alternative to torturing this innocent child was letting the entire city of Boston perish. That fact is simply irrelevant. Torturing an innocent child is *wrong*, just because of what it is in itself, no matter what the consequences. But likewise, letting Boston be annihilated is wrong. It is wrong to allow such massive loss of innocent life when you could have avoided it."

To summarize what we have learned so far: When we examine the features of ethical views that give rise to moral outcome luck and moral dilemmas, we

---

the alternative to B is A, then B is wrong, since it allows a greater evil. By contrast, in a forced choice between B and D: Allow the destruction of New York City, then choosing B is not wrong, under a lesser evil logic.

[25] "War and Massacre", p. 54. See also "The Fragmentation of Value" in *Mortal Questions* (Cambridge: Cambridge University Press, 1979), pp. 128–41.

[26] Walzer, "Political Action: The Problem of Dirty Hands" (op. cit.), p. 164.

[27] Ibid. pp. 167-8.

notice a hidden connection. Proponents of moral dilemmas and of moral outcome luck are united by their commitment to parochial evaluation. Proponents of moral luck are parochialists about outcomes; proponents of moral dilemmas are parochialists about *options*. The question now is: should *we* be either?

It is fair to say that the thought of admitting genuine moral dilemmas and moral outcome luck fills most people – philosophers and lay folk alike – with deep unease. Philosophers have objected to moral dilemmas both on formal[28] and on moral grounds.[29] Perhaps the most popular anti-dilemma argument is grounded in the principle that 'ought' implies 'can': Suppose you're in a situation where you have only two available courses of action, A and B. If this is a moral dilemma, then doing *either* A or B is wrong. You ought to *avoid* doing A and you ought to *avoid* doing B. But according to what is called the Principle of Agglomeration, if a person ought to do one thing and ought to do another, then that person ought to do *both* things. However, avoiding *both* A and B is, by assumption, *impossible*. Hence, this situation cannot be a moral dilemma, on pain of violating the extremely plausible principle that 'ought' implies 'can'.[30]

Moral outcome luck, too, seems to butt up against a powerfully intuitive notion: the idea, labelled by Nagel the "Control Principle", according to which an agent cannot be morally assessable for features of her action that were not under her control.

Yet despite these reasons for unease, many philosophers have thought themselves compelled to accept the reality of moral dilemmas and of moral outcome luck, in order to account for the 'moral phenomenology' of cases like *Ticking Time Bomb* and *Two Drunk Drivers*. In the former case, it seems hard to deny that, however the agent acts, he will have reason to feel guilty. But guilt, it is commonly thought, is appropriate only if one has acted wrongly.[32] And the most compelling evidence for moral outcome luck, it has seemed to many, are

---

[28] See, for instance, W.D. Ross, *Foundations of Ethics* (Oxford: Clarendon Press, 1939), pp. 88-89 and Earl Conee, "Against Moral Dilemmas," *Philosophical Review* (1982), pp. 87-88.

[29] See, for instance, Geoffrey Sayre-McCord ,"A Moral Argument Against Moral Dilemmas" (unpublished manuscript).

[30] An early discussion of this argument is in Williams, "Ethical Consistency" (op. cit.), p. 118.

[32] Versions of this argument are put forward, inter alia, by Bernard Williams, "Ethical Consistency"; Ruth Barcan Marcus, "Moral Dilemmas and Consistency"; Michael Walzer, "Political Action: The Problem of Dirty Hands"; Lisa Tessman, *Moral Failure: On the Impossible Demands of Morality* (New York: Oxford University Press, 2015), especially chapter 2.

the wildly different reactive attitudes that seem appropriate vis-à-vis Ursula and Lara, both for those affected by their actions, as well as for Ursula and Lara towards themselves.

My aim in what follows is to show that we can resist the force of these arguments. With regard to the problem of moral dilemmas, I will argue that there is an alternative way of accounting for the phenomenology of cases like *Ticking Time Bomb*, thereby removing the need to cast them as genuine moral dilemmas. With regard to the problem of moral luck, my strategy will be to defang the problem, by showing that there either is no moral luck or, if there is there is no tension between the Control Principle and the form of moral luck that we have reason to accept.

Key to my argument is a distinction between two perspectives on the evaluation of actions – what I call "the two faces" of morality.

4. The Agent- and Patient-Centric Perspectives

We are both the agents and the patients of actions. We are active in the world, and are in turn affected by the activity of others. Accordingly, we should distinguish two perspectives from which we can morally evaluate a given action: an *agent-* and a *patient-centric* perspective.

From the agent-centric perspective, the central question is: "How well or poorly has the agent comported herself, morally, in φ-ing?" I am using "comported herself" as a term of art: it comprises, first of all, questions about the action's *deontic* status: was it permissible or impermissible , superogatory, etc.? But questions about how well or poorly an agent has comported herself go beyond this; they also include questions about the reasons for which the agent acted, and what excuses she may have had for acting in this way. These are questions that do not typically bear on the permissibility of an agent's action, but can matter to how well or poorly the agent has comported herself.[33]

---

[33] Proponents of the Doctrine of Double Effect (DDE) would dispute that the reasons for which an agent acts do not typically affect the permissibility of her action. For a spirited defense of the doctrine, see Ralph Wedgwood, "Defending Double-Effect", *Ratio* Vol. 24, No. 4 (2011), pp. 384-401. For two prominent critiques, see Judith Jarvis Thomson, "Physician-Assisted Suicide: Two Moral Arguments", *Ethics* Vol. 109, No. 3 (1999), pp. 497–518 and T.M. Scanlon, "The Illusory Appeal of Double-Effect" in his *Moral Dimensions: Permissibility, Meaning, Blame* (Cambridge, MA: Harvard University Press, 2008). Though I am myself on the side of the DDE skeptics, this disagreement will not matter for what follows. Proponents and opponents of the DDE can agree

By the same token, an agent can have comported herself poorly, although her action was permissible, if she did the right thing for very bad reasons. If an agent has comported herself poorly, either because she acted wrongly, or because she did the right thing for very bad reasons, I will say that the agent's conduct has been morally *faulty*.

By contrast, the central question from the *patient*-centric perspective is: "How have *other people* been *affected* by the agent's action? Specifically, has anyone been *wronged* by the agent's behavior? If yes, how seriously?".

I shall present a fuller characterization of what it is to wrong a person in Section 6; but a few preliminary remarks are in order: Whereas "acts wrongly" is a *monadic* predicate, which attaches to an agent (or group of agents) ("X acts wrongly by φ-ing"), wronging is a *dyadic* or '*bi-polar*' moral notion, linking an agent *and a patient*.[34] Claims about wronging have the following form:

X wrongs *Y* by φ-ing.

A wrong is something done by an agent *to a patient*.[35] For Y to complain that she has been wronged by X's φ-ing is not merely – perhaps not even necessarily[36] – to claim that X's conduct was morally wrong or faulty. It is to make a claim about the way X's conduct affected *her.* If X's φ-ing *wrongs* Y, Y has suffered a *moral injury*; she has been mistreated by X. This impairs the relationship between X and Y, and alters it in important ways: X and Y now relate to one another as wrongdoer[37] to victim. As the victim, Y may have grounds to resent X, or to complain based on the injury she suffered, to demand

---

that, if an agent acts for morally awful reasons, she comports herself poorly qua agent – whether or not this implies that her action itself is also impermissible.

[34] For more on the distinction between monadic and bi-polar normativity, see Michael Thompson, "What is it to Wrong Someone? A Puzzle about Justice" in Wallace, R. J., Pettit, P., Scheffler, S., and Smith, M. (eds.), *Reason and Value: Themes from the Moral Philosophy of Joseph Raz,* (Oxford: Clarendon Press, 2004).

[35] If we think there are 'wrongs to self', X and Y can correspond to the same individual. But here, too, wronging retains its dyadic structure. We are talking about one individual in-her-capacity-as-agent and in-her-capacity-as-patient.

[36] See my argument in Section 6.

[37] Or more precisely, but less idiomatically, as "wronging-doer" to victim. What I am interested in is X as someone who *wronged* another not as someone who has acted wrongly.

compensation, or to forgive the transgression. X, in turn, may have reason to apologize, offer restitution, or feel guilty at having committed a wrong.[38]

As we have seen, the question whether there are moral dilemmas or moral outcome luck turns on whether, in determining if an agent has acted wrongly or how blameworthy she is for how she acted, we should endorse parochial evaluation about options or about outcomes. Questions about the wrongness of the agent's action belong squarely to the *agent*-centric dimension of evaluation; and many philosophers have an understanding of what it is to be blameworthy, according to which that question, too, is exclusively a matter of how well or poorly the agent comported himself in φ-ing, and hence are the province of purely agent-centric evaluation as well.[39]

But, as indicated in Section 1, it will be profitable to come at these questions obliquely, by setting aside for the time being the question whether we should endorse parochial evaluation with regard to these *agent-centric* questions, and instead first inquiring into the *patient*-centric perspective. In asking whether, or how seriously, X has wronged Y by φ-ing, should we endorse parochial evaluation?

5. Wronging and Moral Outcome Luck

Whether or not we should accept that an agent's level of *blameworthiness* can vary depending on factors beyond her control, it is hard to deny that there is outcome luck of a kind when it comes to the question how seriously other persons have been *wronged* by the agent's behavior.

It is commonly thought that there is a close connection between a patient being *wronged* and that person having a *right* infringed. While the precise nature of the relationship is contested[40] and will be more closely examined in the following section, it is generally accepted that

> If X *violates* Y's right by φ-ing, X wrongs Y.

---

[38] Since X wronging Y affects not just Y, but also's X's moral relationship with Y, an alternative label for the "patient-centric" perspective could be the "interpersonal" perspective.
[39] We shall return to this issue in Section 9.
[40] For an important recent contribution, according to which a patient can be wronged without having a right that was infringed, see Nico Cornell, "Wrongs, Rights, and Third Parties", *Philosophy & Public Affairs*, Vol. 43, No. 2 (2015), pp. 109-143.

"Violate" is here used in the technical sense first introduced by Judith Thomson. [41] According to this usage,

> X *violates* Y's right by φ-ing if and only if (i) Y has a right that X not φ, (ii) X φs, and (iii) X *acts wrongly* in doing so.

By contrast, I shall say that

> X *overrides* Y's right by φ-ing if and only if (i) Y has a right that X not φ, (ii) X φs, but (iii) X acts *permissibly* in doing so.

Following Thomson, I shall employ the locution "to *infringe* a right" as a neutral expression, which doesn't imply either that a right was violated or merely overridden. Thus, when I override a right, I permissibly infringe it, whereas if I violate a right, I impermissibly infringe it. So understood, the above claim is extremely plausible. If X infringes Y's right by φ-ing while *also* acting impermissibly in so doing, it is surely the case that Y has been wronged.

Let us now consider the *Two Drunk Drivers* case. Unlucky Ursula who kills a child while driving drunk has obviously *violated* a right of that child, namely the child's right against being physically harmed. Thus, Ursula has wronged the child. Moreover, since the right against being physically harmed is a particularly stringent right, the wrong that Ursula has caused the child is correspondingly serious.

By contrast, through sheer luck Lara hasn't violated anyone's right against physical harm. After all, whether a person's right against harm has been violated is a function of whether that person was *actually* harmed, not whether they might have been harmed.

This is not to claim that Lara didn't wrong anyone *tout court*. Arguably, Lara, like Ursula, did wrong various other persons along her route, by exposing them to an unreasonable *risk* of being harmed in an accident. It is plausible that people have a right, not just against being harmed, but against being exposed to such unreasonable *risks* of harm.[42] But these wrongs, for what they're worth,

---

[41] See her *Rights, Restitution and Risk: Essays in Moral Theory,* (Cambridge and London: Harvard University Press, 1986) p. 51.

[42] For a recent defense of this claim, see John Oberdiek, *Imposing Risk: A Normative Framework,* (Oxford: Oxford University Press, 2017).

seem relatively minor compared to the grievous wrong that Ursula inflicted on the child she killed.[43]

Hence, all things considered, it appears that through sheer good luck, Lara didn't end up inflicting nearly as serious a wrong on anyone as did her unlucky counterpart. In determining the degree to which an agent has wronged others, parochialism about outcomes thus seems obviously correct.

By contrast, whether we should also be parochialists about *alternatives* when asking whether X has wronged Y by φ-ing is a more difficult question. To answer it, we need to first step back and ask a more basic question: What exactly is it to wrong a person?

6. What is it to wrong a person?

As I remarked in the previous section, it is generally thought that there is a close connection between a patient being *wronged* and that person having a *right* infringed. For some philosophers, this connection is very simple. They hold:

(1) X wrongs Y by φ-ing if Y has a right that X not φ. [44]

(1) is endorsed by Elizabeth Anscombe, who in fact affirms a stronger *identity* thesis:

A wrong is an infringement of a right. What is wrong about an act that is wrong may be just this, that it is a wrong.[45]

---

[43] This is why Ursula, who imposed a lesser risk of harm on other people, but who was unlucky and ran over a child, has arguably wronged people more even than reckless Rita, who imposed a far graver risk of harm on people, but had a lucky escape and didn't do any actual damage.

[44] Since (1) states only a *sufficient* condition on wronging, it is compatible with the thought that there may be ways of wronging a person other than by overriding one of her rights. This has recently been suggested by Nico Cornell, "Rights, Wrongs, and Third Parties". I will bracket this question in what follows, as it will not bear on my argument.

[45] G.E.M. Anscombe, "On the Source of the Authority of the State," in *Authority*, ed. Joseph Raz (New York: New York University Press, 1990), p. 152.

If (1) were correct, it would follow that X can wrong Y even while acting permissibly, all things considered. After all, there can be cases where Y has a right that X not φ, yet X permissibly infringes that right.

On reflection, however, (1) appears too permissive. There are cases where an agent infringes a patient's rights, yet it seems the agent hasn't wronged the patient. The following well-known case from Joel Feinberg provides an illustration[46]:

> *Feinberg's Cabin:* You are hiking in the mountains when a sudden winter storm strikes with such force that your life is endangered. Through a stroke of luck, you stumble across a cabin, clearly someone else's private property (Joel Feinberg's, as it happens!), but currently unoccupied. You smash a window and gain entry. Over the next three days, to keep from freezing, you burn some of Feinberg's wooden furniture.

Clearly, you have not acted impermissibly by breaking into Feinberg's cabin and burning his furniture. Private property is not sacrosanct. At the same time, it does seem that in the course of doing what you needed to survive, you *infringed* various property rights of Feinberg's . It is not that, because you have a justification of necessity for breaking into his cabin and burning his furniture, Feinberg *has* no property right in his cabin and his furniture in these circumstances. If that were the case, it would be hard to account for the fact that, even though it was permissible for you to burn Feinberg's furniture, once the storm has passed you have a duty to compensate him for its loss.[47]

But even though Feinberg retained a property right in his furniture, which you infringed by burning it, is it really apposite to say that you 'wronged' Feinberg by burning his furniture? My sense is that this would be infelicitous.

Why is that? What more is it for X to wrong Y by φ-ing than for X to infringe Y's right that X not φ?

I think the missing ingredient is this: There is a conceptual connection between wronging and *relational impairment*. To wrong a person, I submit, is to cause them a moral injury which impairs your relationship with them. This

---

[46] After Joel Feinberg, "Voluntary Euthanasia and the Inalienable Right to Life", *Philosophy & Public Affairs*, Vol. 7, No. 2, (1978), pp. 93-123.

[47] For an elaboration of this point, see Judith Thomson, "Some Ruminations on Rights" and "Rights and Compensation", both in her *Rights, Restitution, and Risk: Essays in Moral Theory.*

impairment can take various forms. If X has wronged Y, Y may have grounds to feel moral emotions like resentment or anger. Y may also have reason to alter her standing dispositions and attitudes towards X: she may reduce her readiness to interact with X in the future, or to be his friend; she may withdraw her attitude of goodwill towards X, ceasing to hope that things will go well for X and no longer taking pleasure when they do; and so on. [48] For his part, having wronged Y gives X reason, not just to *compensate* Y for whatever material loss she may have suffered, but to try to undertake a project of *relational repair*. Wronging a person, as opposed to merely overriding her rights, makes it fitting to seek *forgiveness*, with the aim of achieving *reconciliation*.[49]

It might be objected that we surely can wrong people with whom we have no relationship, and hence with whom there is no relationship to be impaired. But the sense of relationship that I have in mind, and which I claim is impaired by wronging a person, is not a "thick" interpersonal relationship like friendship (which I have only with certain other people) but what Tim Scanlon calls the default "moral relationship" which I have with *all* other people simply in virtue of the fact that they are moral agents. It is this default moral relationship which, I claim, is impaired by wronging a person. For instance, I think that, even towards complete strangers, with whom I have no thick interpersonal relationship to speak of, a basic attitude of *good will* is called for. I have reason

---

[48] This account of relational impairment draws heavily on the much fuller account offered by T.M. Scanlon in Chapter 4 of *Moral Dimensions: Permissibility, Meaning, Blame*, (Cambridge, MA: Harvard University Press, 2008).

[49] To deny that X wronged Y by φ-ing is not to deny that X's φ-ing may have produced a *moral remainder*, as witnessed by the fact that X may have a duty to compensate Y for having φ-ed. More than that, it may be appropriate for X to *apologize* to Y for having φ-ed. But here it is important to draw a distinction between two types of apologies. In some contexts, apologies function as pleas for *forgiveness*. To apologize to a person in this sense is to acknowledge that your conduct has impaired your relationship, and to seek to repair this damage. In other contexts, apologies function as *explanations* of your behavior. The aim of this type of apology is not to repair a damaged relationship; it is precisely to explain that no such impairment has occurred. Suppose I miss my lunch appointment with you. Later you receive this note: "Apologies for missing lunch. My mother had a heart attack and I had to drive her to hospital." My apology here explains why, given my reasons for missing lunch, I did *not* wrong you. Moreover, as Julia Driver points out, "if [you] were to later say "I forgive you," that would, charitably, be taken as a joke." (See her "Wronging, Blame, and Forgiveness" in David Shoemaker (ed.) *Oxford Studies in Agency and Responsibility* (Oxford: Oxford University Press, 2017)). Hence, if X did not wrong Y by φ-ing, this doesn't rule out that X may have reason to offer Y an apology in the sense of giving Y an explanation for his behavior. But an apology in the sense of asking for forgiveness would be uncalled for.

to hope that things go well for these strangers, to take pleasure when they do, etc. But this basic attitude of good will, which is part of the default moral relationship, might permissibly be *withdrawn* if the stranger wrongs me. Likewise, certain reactive attitudes, such as resentment and anger, which would have been inappropriate in the absence of wronging, may now be appropriate. This is what I will have in mind when I say that wronging is connected to relational impairment. It is the impairment of the basic moral relationship that we have to every other person, simply in virtue of their being a moral agent. (This of course is entirely compatible that by wronging another person I *also* impair various thicker relationships, such as friendship, that I may stand in with regard to this particular person)/

If these observations are on the right track, then (1) should be amended to

(2) X wrongs Y by φ-ing if (i) Y has a right that X not φ, and (ii) X's infringing of this right causes Y a moral injury that impairs X's relationship with Y.

I don't think that there is such relational impairment in *Feinberg's Cabin*. Burning Feinberg's furniture to survive the blizzard calls for compensation, to be sure; but asking him for 'forgiveness' seems misplaced. [50]

Feinberg's Cabin stands in stark contrast to cases like the following:

---

[50] On the other hand, suppose that, having burned Feinberg's furniture to survive, you later *refuse* to compensate him, although you are able to. That, unlike the initial burning of the furniture, *would* impair your relationship. And in that case, it seems right to say that Feinberg has been wronged.

  Likewise, suppose that prior to burning Feinberg's furniture, you had the ability to contact him to seek his permission (his telephone number is on a post-it note in the cabin), yet you failed to do so. Here, again, I think you may have wronged Feinberg, by failing to seek his permission. This is so, even though I also believe that, had you asked Feinberg and had he *refused* to grant his permission, necessity would have permitted you to burn his furniture nonetheless, and burning his furniture would not have wronged him. (Incidentally, this is exactly the view of Hugo Grotius in *The Rights of War and Peace* (1625), Book II, Chapter 2, Sections 6-9). This combination of claims may strike us as puzzling. If you do not need Feinberg's permission in order to permissibly burn his furniture, then why do you wrong him if you do not seek to obtain his permission first? Part of the explanation may be that seeking Feinberg's permission is the appropriate way of showing respect for his property rights. This claim is structurally analogous to some theories of civil disobedience (such as Rawls's), according to which, while civil disobedience may be permissible in response to serious injustices, *fidelity* to the law requires citizens to have first exhausted other, lawful means of political reform or resistance.

> **Large Man:** A runaway trolley is headed for five people trapped
> ahead on the tracks, whom it will run over and kill if not
> stopped. You are standing on a bridge overlooking the tracks. If
> and only if you push a large man standing next to you onto the
> tracks ahead of trolley, the trolley will be stopped and the five
> will be saved. The large man, however, will be crushed by the
> trolley and die. You push the large man onto the tracks.

What you do to the large man seems a *paradigm* of wronging a person. And
indeed, infringing the large man's right to bodily inviolateness by pushing him
to his death in this manner clearly impairs your relationship, in the various
senses detailed above. (*Killing* a person, in a way that wrongs him, is an extreme
form of relational impairment).

Since most (non-consequentialist) philosophers would judge that, in
addition to wronging the large man, you also acted *impermissibly* in this case,
this might suggest the following account of when infringing a person's right
leads to relational impairment, and thereby wrongs them:

> (3) X wrongs Y by infringing one of her rights if and only if X
>   *violates* Y's right, i.e. infringes it *impermissibly*.

This is exactly the account of wronging endorsed by Judith Thomson[51] and it is
close to orthodoxy among philosophers thinking about these issues. If this
orthodox account of wronging were correct, the crucial difference between
Feinberg's Cabin and *Large Man* is that in the former case, you infringe
Feinberg's right permissibly, thus merely *overriding* it, whereas in latter case,
you *violate* the large man's right, because your action is all-things-considered
impermissible.

However, I believe that orthodoxy has got it wrong. If to wrong a person is
to cause her a moral injury that impairs your relationship with her, then
Thomson's necessary condition (you have wronged another person only if your
action was *impermissible*, all things considered) seems too stringent. To see this,
consider

---

[51] See Judith Jarvis Thomson, *The Realm of Rights* (Cambridge, Mass.: Harvard University Press, 1990), p. 122.

Large Man II: You push a large man onto the tracks in order to save *n* people from an onrushing trolley. The large man is killed by the trolley.

Most of us are not moral absolutists. Rather, we think, as threshold deontologists do, that if the number of lives we could save by sacrificing the large man becomes sufficiently huge, we may permissibly infringe his rights and push him, for the prevention of catastrophic moral horror. Suppose that *n*, the number of people actually on the tracks, was just above this critical threshold. In that case, you have acted permissibly, according to this threshold deontological picture.

Thomson's account then implies that, since your action was permissible, you cannot have *wronged* the man by pushing him to his death. But that seems hard to believe. Considerations of the greater good may block the charge that you acted *impermissibly*; they do not block the charge that, in so doing, you *mistreated* the large man, in a way that impaired your relationship. It would be grotesque to suggest that, since you haven't acted impermissibly, it would therefore be misplaced to ask the man (or his relatives) for forgiveness.

The key mistake of the orthodox analysis is this: The distinction between

(a) X *violates* Y's right by φ-ing, and
(b) X *overrides* Y's right by φ-ing

is a function of an <u>agent-centric</u> question: 'Did X act all-things-considered permissibly in φ-ing?'. By contrast, the distinction between

(c) X *wrongs* Y by φ-ing, and
(d) X *does not wrong* Y by φ-ing

is a <u>patient-centric</u> distinction, which tracks whether Y's relationship with X was impaired by X's φ-ing. But there is no reason to suppose that that which makes a crucial difference to the first, agent-centric question – 'How many lives did X save by pushing Y in front of the trolley, and was that number above the critical threshold?' – necessarily holds the same relevance for the patient-centric question.[52] There is thus no reason to suppose that the line which separates (c)

---

[52] Imagine Y (or his ghost), upon learning that the number of lives that X saved by killing him was indeed above the relevant threshold, saying bitterly to X: "Good for you; at least you didn't

from (d) is drawn in the same place, and on the basis of the same considerations, as the line which separates (a) from (b).

I will now present what seems to me a better account of when infringing a person's right leads to relational impairment, and thereby wrongs them. Return to this question: Why is Feinberg not wronged when you burn his furniture to survive, whereas the large man *is* wronged when you push him onto the tracks to save five lives?

The crucial difference, I believe, is this: Though in both cases the agent infringes a right of the patient, in the former the agent acts in a way that the patient could not, in the circumstances, have permissibly *denied his assent to*. I shall understand this notion as follows:

> Y cannot permissibly deny her assent to X's φ-ing in circumstances C just in case, were Y able to decide whether X φs in circumstances C or not, it would be impermissible for Y to make it the case that X does not φ.[53]

Thus, to say that, in the midst of the blizzard Feinberg cannot permissibly deny his assent to your breaking into his cabin and burning his furniture is to say that, even if, hypothetically, he could stop you from doing these things by refusing his go-ahead, it would be morally impermissible to do so. Though breaking into Feinberg's cabin and burning his furniture infringes his property rights, these are rights that Feinberg couldn't permissibly *exercise* in the circumstances. There is thus a kind of "moral unanimity" between you and Feinberg: You could truthfully say to him: "I am not doing anything to you (or your property) that you, if you were fully moral, would not assent to".

Things are quite different vis-à-vis the large man. The large man *can* permissibly deny his assent to being pushed in front of the trolley, to save five lives. Had it been up to him to decide whether you push him off the bridge, he

---

act *impermissibly*. But don't think for a moment that this makes things alright between *you and me*."

[53] Suppose that X says truthfully to Y: "I will only φ if you give me the go-ahead." For it to be permissible for Y to deny her assent to X's φ-ing is for it to be permissible for Y not to give her go-ahead to X's φ-ing. Note that this is weaker than, and does not entail, the claim that Y can permissibly take whatever means are necessary to prevent X from φ-ing.

could permissibly have refused to give his go-ahead. Hence, pushing the man off the bridge opens up a gap between how you *actually* treated him and the kind of treatment he was morally required to assent to. There is no moral unanimity between you and him. In a sense, you are just *overpowering* him. You are subjecting him to treatment which he wouldn't be required to endure, were this up to him. What is stopping him from enjoying his right to physical inviolateness is not a *moral* requirement not to exercise his right and allow it to be infringed, but sheer force. This gives the infringement of his right a normative significance which the infringement of Feinberg's right does not possess. And that makes for a crucial difference between the cases. It plausibly explains why the behavior your inflict on the large man constitutes a moral injury and damages your relationship, whereas what you do to Feinberg only requires compensation but does not constitute a wronging.

If this is correct, we can adopt the following precisification of our relational account of wronging:

(4) X wrongs Y by φ-ing if (i) Y has a right that X not φ, and (ii) Y could have permissibly (and rationally[54]) denied her assent to X's φ-ing.[55]

My account of wronging does not just capture our judgments about *Feinberg's Cabin* and *Large Man*. Unlike the orthodox account, it also supports the intuitively correct analysis of *Large Man II*, namely that here too, if you push

---

[54] This qualification is needed to deal with a wrinkle which arises in a case from Tom Dougherty, "Why Does Duress Undermine Consent?" (forthcoming in *Nous*) Consider:

> **Kicking the colleague:** "D is about to fall asleep in an important meeting, and will be hauled over the coals by her boss if she does. E is sitting across from D, and notices this. She judges that the only way to prevent D falling asleep is kicking D under the table, and so she kicks D."

Since D has a right against being kicked by others, E thereby infringes D's right. But E does not thereby *wrong* D. Although being kicked is something that D could *permissibly* deny her assent to, it wouldn't be *rational* for her to deny her assent, because in this case it is clearly for her own greater good.

[55] Victor Tadros embraces something close to the inverse of (4). According to Tadros, Y is liable to be harmed by X at some cost, c, for the sake of some goal, g, if Y is required to bear c for g, or would be required were Y able to bear c for g, and if Y is liable to bear c for g, Y is not wronged by inflicting c on her for g. See his "Duty and Liability", *Utilitas,* Vol. 24, Issue 4 (2012), pp. 259-277 and "Causation, Culpability and Liability" in Coons and Weber (eds.), *The Ethics of Self Defence* (New York: Oxford University Press, 2016), pp. 110-130.

the large man off the bridge, you wrong him, despite the fact that your action is morally permissible.

How so? The crucial point is this: The fact that in *Large Man II* it is permissible to push the man off the bridge does *not* make it the case that the large man is morally required to assent to this treatment. In general

X acts permissibly in infringing Y's right that X not φ

does not imply that

Y is morally required to give his assent to X's φ-ing.

Frances Kamm provides a particularly clear illustration. Suppose we are in the standard Trolley Problem: X is the bystander, about to pull a switch that will redirect the trolley away from five persons trapped on the tracks ahead and towards Y, who is trapped on a side-track, saving the five but killing Y. Would Y have to assent to this? Intuitively not. Indeed, Kamm makes a stronger claim: suppose that Y actually has the ability to physically prevent X from turning the trolley onto him; it would not be wrong for Y to exercise it, she believes:

> [I]t would not be impermissible for the one person toward whom the trolley is redirected to resist our doing this. For example, if he could press a button and send the trolley away, back where it came from, even if we or the five originally threatened would be killed, I think this would be permissible.[56]

What an agent can permissibly do to a patient in pursuit of the greater good, and what the kinds of treatment the patient is required to assent to, in light of the same greater good considerations, can come apart.

---

[56] Frances Kamm, *Intricate Ethics*, (New York: Oxford University Press, 2007), pp. 231-2. Like myself, Kamm also thinks that "[t]he permissibility of someone's resisting our permissible act is (…) evidence for the fact that we still wrong him in acting permissibly and in (permissibly) infringing his right." (ibid.) See also her *The Trolley Problem Mysteries* (New York: Oxford University Press, 2016). I have given a fuller explanation of why this is the case. Another philosopher who makes the case (albeit on different grounds from those advanced here) that there can be permissible actions that nonetheless wrong a person is Julia Driver in "Wronging, Blame, and Forgiveness" (op. cit.).

There are two general reasons for this: First, as many philosophers have argued, individuals have an agent-centered prerogative to be partial to their own interests (by, as it were, applying a "multiplier" to them in practical deliberation).[57] As a result, the sacrifice that X can permissibly impose on Y, in order to bring about a greater good of size G, is larger than the sacrifice that Y must assent to having imposed on her by X for the sake of bringing about G. As we saw above, X can permissibly turn the trolley onto Y to save five lives. But, if it were up to Y, she could permissibly refuse to have this sacrifice imposed on her. This agent-centered prerogative is not limitless, however. If the sacrifice Y had to endure was much smaller – if she had only to lose a little finger to save five lives – Y couldn't permissibly refuse her assent. And, of course, as the greater good considerations get weightier and weightier, the severity of the sacrifices that a person must assent to becomes larger in turn.

However, second: While there may be no upper limit to the burdens that an agent may permissibly impose on a patient for the sake of a sufficiently important greater good, I believe that there *is* an upper limit to the sacrifices that Y can be required to assent to, no matter how large the greater good. Y would be a saint or hero if she gave her assent to sacrifices beyond this limit. But doing so is supererogatory.[58] Call this the Upper Limit Claim.

I think that a combination of these two factors explains the intuitive datum that in *Large Man II* you wrong the victim by permissibly pushing him in front of the trolley. While we stipulated that the number of lives you save is great enough that sacrificing the large man is permissible, he isn't required to assent to this sacrifice. This is so, both because he permissibly gives greater weight to his own interests in practical deliberation than you do, and arguably also because the sacrifice imposed of him – to give up his *life* for the greater good – is more than can be morally *required* of any person, whatever the stakes.

---

[57] The locus classicus is Samuel Scheffler, *The Rejection of Consequentialism*, (Oxford: Clarendon Press, 1982).

[58] This is not to exclude the possibility that, in some cases, refusing to assent to a supererogatory sacrifice, while permissible, could nonetheless expose a person to justified moral criticism. There may be what Elizabeth Harman calls "morally permissible moral mistakes". See her "Morally Permissible Moral Mistakes", *Ethics*, Vol. 126, No. 2 (2016), pp. 366-393. For instance, an agent who refuses to sacrifice herself, even to save thousands of others, may reveal a level of self-regard that is not morally decent. However, the fact that she can be criticized for the way in which they chose to exercise their personal prerogative does not imply that her decision was morally impermissible.

The permissible partiality that persons have towards themselves and the Upper Limit Claim thus account for one of the ways in which the agent- and the patient-centric evaluations of actions can come apart. The patient can permissibly deny her assent to some sacrifices that the agent can permissibly impose on her. That is why X, while acting permissibly, can nonetheless *wrong* Y.

7. Purported moral dilemmas and wronging

My account of wronging holds important lessons for how to think about the purported cases of moral dilemmas from Section 3. It suggests that – even if we reject the characterization of these cases as genuine moral dilemmas, where no matter how the agent acts, his action is impermissible – the following weaker claim is nonetheless very plausible: These are situations where the agent cannot avoid *either* acting impermissibly *or else* wronging a person.

Suppose Walzer is right that the "correct" course of action in *Ticking Time Bomb* is to torture. The destruction of an entire city is such a catastrophic outcome that allowing it would be impermissible. But assume we deny what Walzer wants to also assert, namely that torturing the innocent child is *nonetheless* morally wrong, all things considered. Under the orthodox account of wronging, this would prevent us from saying that, in torturing the child, you do her a grievous wrong. For, according to orthodoxy, you wrong a person only if you behave impermissibly.

Not so on the account of wronging that I have defended. On my view, you do wrong the child because you impose on her a sacrifice – undergoing horrific torture – that she arguably is not morally required to assent to. The *Ticking Time Bomb* scenario thus confronts you with a forced choice between either acting wrongly or grievously wronging a person. In *this*, weaker sense, there truly is no way of escaping this situation with clean hands.

Mutatis mutandis for *Sartre's Student.* The moral urgency of fighting the Nazi invader was very great in 1940s France. Let us assume that Sartre's student is correct in thinking that joining the resistance is his moral duty, which it would be wrong to shirk. So he goes off to fight the Nazis, abandoning his mother. Again, even if we refuse the dilemmatic claim that doing so was also impermissible all things considered, it is very plausible that his decision had a moral remainder: he wronged his mother by deserting her. Parents, especially

ailing parents, have a right to the care and attention of their children, and I do not think it would have been impermissible selfishness on his mother's part to refuse to assent to her own abandonment. Hence, even if the student acted permissibly there was a moral price to pay: he could do the right thing only by wronging his mother.

There is a second lesson. If the Upper Limit Claim is correct, then for some tokens of certain act-types (killing, torture, rape, etc.) it is the case that if an agent subjects an innocent person to such treatment, he has *wronged* the person, *no matter* what the alternatives, however terrible.[59] For these are forms of treatment that a patient is not required to assent to under any circumstances. The modal context of available alternatives is irrelevant.[60] Thus, while there are reasons to doubt parochialism about options as a thesis about the agent-centric notion of *wrongness* – it is prone to generate dilemmas – we should sometimes be parochialists about the patient-centric question whether a person has been *wronged*.

8. Tragic moral conflicts and 'directed' guilt

Once we come to think of cases like *Ticking Time Bomb* and *Sartre's Student* as situations where all the agent's available courses of action would involve either acting wrongly or else seriously wronging some person, the phenomenological case for moral dilemmas loses much of its force.

---

[59] The claim must be stated in terms of act-tokens rather than act-types for the following reason: Do I wrong you by killing you? In some situations, this is a non-parochial matter. Suppose you are howling in agony, and so communication is out of the question. If I fail to act, you will suffer this unbearable agony for many months before dying. I also have the option of involuntarily euthanising you, as a mercy-killing. I think I would not wrong you by killing you in this situation. This shows that it is not the case that for the act-type killing, whether you are wronged by a killing is non-parochial. What remains is the possibility that for some act-tokens of killings, e.g. the token killing in the *Large Man* and the *Large Man II* case, whether the victim is wronged is non-parochial. I thank Tom Dougherty for discussion of this point.

[60] This will not, of course, be true for *all* cases where a person is wronged: Often, whether a person was wronged does depend on what the agent could have done instead: For instance, it would wrong me if you caused me to lose my pinky just in order to save someone else's pinky. But it wouldn't wrong me if you cause me to lose my pinky in order to save the lives of 100 people, whom you would otherwise have had to let die. The modal context will be irrelevant only at the "top end of the scale" of sacrifices, to which the Upper Limit Claim applies. But it is exactly in this region where many (putative) moral dilemmas are at home. They traffic in extreme situations.

Proponents of moral dilemmas such as Williams, Marcus, Walzer, or Tessman appeal to the characteristic phenomenology of 'hard cases' or 'tragic moral conflicts' – the feeling that, however the agent resolves the conflict, he ought to feel guilty for what he does – to argue that the agent must be facing a genuine moral dilemma. Guilt, Walzer stresses, is distinct from feeling *badly* about what one has done, nor is it to be confused with feeling *melancholy* or *sadness* or *regret* at having been forced, by dire circumstances, to do terrible things. Latter emotions are all compatible with recognizing that one has acted rightly and blamelessly. By contrast, it is appropriate for an agent to feel *guilty* only if only if she properly believes that she has acted wrongly.[61]

This phenomenological case for moral dilemmas relies on what I call an "orthodox" view of guilt (echoing the orthodox view of wronging discussed in Section 6), according to which it is appropriate for a person to feel guilty for how she has acted only if she properly believes that she has acted wrongly.

I believe that, once we recognize that wronging a person and acting wrongly can come apart – that we can wrong a person even while acting permissibly – there is little to sustain this orthodox view of guilt. It is a mistake to think that feeling guilty for how one has behaved can be appropriate only when one has acted wrongly. Guilt can also be appropriate – or rather there is a *kind* of guilt which is appropriate – when a person recognizes that he has wronged another, quite independently of whether doing so was also all-things-considered wrong or not.

We should distinguish two types of guilt: "Monadic" guilt about how one has acted is the emotional response made appropriate by the recognition that one has acted wrongly, or otherwise been at fault as an agent (for instance, by doing the right thing, but for awful reasons). This emotion is not directed at another person. Indeed, an agent can feel guilty in this sense even though her action had no adverse effect whatsoever on other people.

By contrast, "directed" or "bipolar" guilt, as in "feeling guilt *towards* someone", is the appropriate emotional response to having *wronged* a person. Feeling guilt *towards* a person indicates that you think your relationship with this person is in need of moral repair. Your conduct has caused them a moral injury and damaged your relationship. You must now seek forgiveness, make amends, etc.[62]

---

[61] "Political Action: The Problem of Dirty Hands", pp. 166-7.

[62] Other languages are clearer than English in acknowledging that emotions of guilt often have a bipolar character. The German word *Schuld*, for instance, translates as both *guilt* and *debt*,

I think we are all familiar with this feeling. Indeed, given the press of modern life and the many competing demands on our time, attention, and energy, it is almost inescapable. In my rush to complete this paper by a deadline, I have neglected relatives, let my correspondence slide, even failed to fully mourn the passing of a dear friend. I feel guilty about these things, though not because I think that my priorities were necessarily unjustifiable or my behavior impermissible. I am pained nonetheless, because I know that my actions had a cost on others and my relationships with them, and because I realize that they would not be wrong to feel aggrieved.

In more serious cases, where the agent has caused others grave moral injury, the attendant feelings of directed guilt can be severe. A politician who orders the torture of an innocent child to save his city from destruction will likely be haunted for the rest of his life. That one's conduct has been faultless is no defense against such feelings. Indeed, *failing* to experience directed guilt, even when one is convinced that one's decision was morally right, can be a sign of serious ethical deficiency. Harry Truman, who bragged that he "never lost any sleep" over having ordered atomic bombs to be dropped on Hiroshima and Nagasaki, strikes us as repellent.[63] Even granting what is now frequently challenged, namely that Truman's decision was necessary to avert an even worse slaughter, the fact that the attacks may have been permissible is *hardly* the only thing that matters.[64] That the atomic bombs killed and maimed hundreds of thousands of civilians who were not liable to attack (and were therefore wronged by being targeted) has a significance quite independent of the question whether the attack was morally justified all things considered. To

---

suggesting a link between guilt and the paradigmatically bipolar notion of indebtedness (you are indebted *to someone*). And unlike "to feel guilty towards someone", which is a somewhat unnatural construction in English, "sich jemandem gegenüber schuldig fühlen" is a common and entirely idiomatic expression in German. Likewise for "se sentir coupable envers quelqu'un" in French.

[63] Sadao Asada, "The Mushroom Cloud and National Psyches: Japanese and American Perceptions of the Atomic Bomb Decision, 1945-1995," in Laura Hein and Mark Selden, eds., *Living With the Bomb: American and Japanese Cultural Conflicts in the Nuclear Age* (New York: M.E. Sharpe, 1997), p. 179.

[64] Michael Walzer himself thinks that the atomic attacks on Hiroshima and Nagasaki could not have been justified under a lesser evil logic and were simply wrong. By contrast, he believes that the Allied terror bombing of German cities in the early stages of World War II was justifiable by appeal to "supreme emergency", though still 'severely immoral'. See Michael Walzer, *Just and Unjust Wars Just and Unjust Wars: A Moral Argument with Historical Illustrations (New York: Basic Books*, 1977), pp. 267-68.

a more decent person, that fact, by itself, would have provided ample grounds for tormenting feelings of guilt.[65]

We can thus vindicate the pre-theoretical sense that in cases of tragic moral conflict the agent finds himself in a situation where he is "damned if he does, damned if he doesn't". There truly is no way for the agent to escape the situation "morally unscathed", without either acting wrongly or else seriously wronging some person. And, however he chooses, he will therefore have reason to feel guilty for what he has done – either monadic guilt, at having acted wrongly, or directed guilt, at having caused others serious moral injury (or in some instances *both*). In that sense, such cases are, on my account as well, truly "tragic". Indeed, the epithet may be an even better fit, since on the story I have told the agent can have reason to feel guilty, in spite of his conduct having been *faultless*.

The mistake of moral dilemma theorists was to assume that we can capture this tragic aspect of hard moral cases only by making the paradoxical claim that, however the agent acts, his action will be wrong. What I have shown is that we can have tragedy without the price of paradox.

9. Blame and the Puzzle of Moral Outcome Luck

Let us return, finally, to moral outcome luck. Here too, the distinction between agent- and patient- centric modes of evaluation can bring much clarity.

The problem of moral outcome luck, as posed by Nagel, can be stated in terms of two propositions, each of which may seem individually compelling but which appear in tension with one another:

(1) The unlucky drunk driver is more blameworthy than the lucky driver, because of how her action turned out.

(2) *The Control Principle*: An agent cannot be morally assessable for features of her action that were not under her control.

---

[65] I remain agnostic on the question whether failing to feel directed guilt when it is appropriate does not just indicate, but *constitutes*, a moral deficiency. Perhaps what is morally required is only that the agent *recognize* that he wronged another. However, in most people such a recognition is typically accompanied by the moral emotion of guilt, and thus failing to experience the emotion is at least *evidence* of a moral shortcoming, if not directly constitutive of it.

The Control Principle is powerfully attractive in its own right. As Nagel writes, "[p]rior to reflection it is intuitively plausible that people cannot be morally assessed (…) for what is due to factors beyond their control."[66] And claim (1), is supported by the phenomenological argument: we cannot deny (1), the argument goes, lest we lose our ability to make sense of very different reactive attitudes that seem appropriate towards Ursula and Lara. The philosophical puzzle of moral outcome luck is how to resolve the apparent tension between (1) and (2).

Before we can sensibly approach this task, we must first become clearer about the notion of blame itself. What is it to blame a person and what is it to be blameworthy? In so doing, however, we are straightaway confronted with the enormous proliferation of accounts in the literature. My aim here cannot be to arbitrate between these manifold accounts. Fortunately, doing so will not be necessary for our purposes. Instead, it will suffice to divide these accounts into two broad clusters:

In the first cluster are accounts of blame and blameworthiness according to which an agent's level of blameworthiness for φ-ing is *solely* a function of the agent's level of fault in φ-ing, i.e. of how poorly she has comported herself *qua agent*.

Susan Wolf, for instance, holds that "blameworthiness is solely a function of faultiness. In other words, equal fault deserves equal blame."[67]

Michael Zimmerman characterizes his conception of blameworthiness as follows: "when I say that a person is blameworthy, I shall mean that her moral record is adversely affected by some (…) fact [about the person]. To praise or blame someone, in this sense, is simply to make a judgment about her moral record, a judgment which may form the basis of, but which is not itself, a "reaction" either in attitude or in some more robust form of behavior toward that person."[68] Zimmerman then goes on to argue that a person's degree of blameworthiness for φ-ing is solely a function of her degree of faultiness.

Finally, on Judith Thomson account, a person's degree of blameworthiness for φ-ing depends on the strength of the reason that her φ-ing gives you for

---

[66] His "Moral Luck" in D. Statman (ed.) *Moral Luck* (Albany: State University of New York Press, 1993), p. 58.

[67] Susan Wolf, "The Moral of Moral Luck", *Philosophic Exchange*, 31 (2001), pp 4–19 at p. 16.

[68] Michael Zimmerman, "Taking Luck Seriously", *The Journal of Philosophy*, Vol. 99 (2002), pp. 553–576, at page 556.

thinking that she is a *bad person*.[69] Thomson stresses that, for two agents whose conduct was equally faulty and whose actions differ merely in terms of their outcomes as determined by luck, you have no stronger reason to think one a bad person than the other.

If we understand blameworthiness as exclusively a function of the agent's faultiness, then it does indeed seem that there couldn't be such a thing as moral outcome luck. For an agent's *faultiness* in φ-ing cannot plausibly be said to vary depending on factors beyond her control: Ursula and Lara's conduct was equally reckless; they imposed an identical level of risk on others; they manifested the same indifference to their safety. *These* are the factors that matter to their level of faultiness, not the outcome of their action, which was beyond their control. Thus, on these accounts of blameworthiness, we must reject (1).

Notice, however, that if this is how we understand what it is for an agent to be blameworthy, then there is in fact *no pressure* to accept (1) on account of the "phenomenological" differences between the lucky and the unlucky driver. For, if blame is solely a matter of *agent-assessment*, then there is no analytic connection between an agent's level of blameworthiness and the reactive attitudes that it is appropriate to have towards her. Having a reactive attitude like blame or resentment is not what blaming the agent *consists* in, on these views; it is only, in David Enoch and Andreï Marmor's helpful phrase, a "*blame-related* reaction", which can be *justified* or *made appropriate* by facts about the agent's blameworthiness.[70]

But while faultiness clearly matters to the reactive attitudes that it is appropriate to have towards an agent, it is not the *only* thing that matters. The appropriate reactive attitudes can also be a function of *patient*-centric facts about how your, or others you care about, were *affected* by the agent's action. Indeed, what you resent a person *for* is not only, perhaps not even primarily, the faultiness of her agency, or the deficient quality of her will, but the unjustified *injury* she caused you, or others you care about. That Ursula inflicts a far graver wrong on the child she runs over than Lara inflicts on anyone, therefore seems undeniably relevant to the reasons that this child, or her family, have for resenting Ursula. While it is fitting to resent both Lara and Ursula for

---

[69] Judith Jarvis Thomson, "Morality and Bad Luck", in D. Statman (ed.) *Moral Luck* (Albany: State University of New York Press, 1993).

[70] David Enoch and Andrei Marmor, "The Case Against Moral Luck", *Law and Philosophy*, Vol. 26, No. 4, (2007), pp. 405-436.

recklessly putting other people at risk, Ursula is the one who *kills* the child, and this is additional ground for resentment.[71]

So, on this first, purely *agent*-centric, way of understanding blameworthiness, we solve the puzzle by retaining the Control Principle and giving up the claim that there is moral luck with regard to blameworthiness. But saying this, I have argued, does not bring us into tension with the phenomenology of cases like *Drunk Drivers*, since that can be fully explained as a function of *patient*-centric differences in the way the lucky and the unlucky drivers *affect* other people.

Consider next the second cluster of philosophical accounts of blame. What these accounts have in common is that, in asking about the agent's level of blameworthiness for φ-ing, we are *not* just asking (either directly or indirectly) about how *faulty* the agent was in φ-ing. Consider two representative examples.

On the classic Strawsonian account, to be blameworthy just is to be the appropriate object of one or more of the reactive attitudes, such as resentment,

---

[71] This is also the view of Michael Otsuka in "Moral Luck: Optional, not Brute", *Philosophical Perspectives*, Vol. 23, Ethics (2009), pp. 373-388. By contrast, in maintaining that the wrongful injury which Ursula causes is independent grounds for resentment, Otsuka and I may be at odds with Peter Strawson himself, who maintains that resentment and indignation are reactive attitudes "to the qualities of others' wills" (P.F. Strawson, "Freedom and Resentment", *Proceedings of the British Academy* 48 (1962), pp. 187-211, at p. 199). Presumably, the quality of a person's will does not vary depending on whether, as luck would have it, she kills or injures someone. On the other hand, Strawson also holds that the strength of our resentment and indignation "is in general proportioned to what is felt to be the magnitude of the injury and to the degree to which the agent's will is identified with, or indifferent to, it" (p. 207). Whatever the truth about Strawson's actual view, an account of resentment that limits the grounds for resentment to the quality of other's wills alone seems untrue to the facts. It can account neither for the very different reactive attitudes that people actually experience in cases like *Drunk Drivers*, depending on what outcome actually results; nor can it make sense of our normative intuitions about the reactive attitudes that seem *appropriate* in such circumstances. To reinforce the latter point, imagine that two sets of parents – the first, the parents of Ute, the child run over by Ursula; the second, the parents of Richard, a child who was put at risk, but not actually injured, by Rita, our *grossly reckless* drunk driver. If it were the case that justified resentment was just a function of quality of will of the agent, then it would be appropriate for Richard's parents to resent Rita more than Ute's parents resent Ursula. For Rita, by going on the road in a far worse state of incapacitation than Ursula, if anything manifested an even greater level of indifference towards other people's safety, and hence a more objectionable quality of will. But it seems quite wrong to suggest that Richard's parents, who didn't actually suffer a bereavement, have greater grounds for resentment than Ute's, who lost their child.

indignation, etc.[72] But, as I just argued, there are compelling reasons to think that what makes an agent an appropriate object of resentment, indignation, etc., are *not* just questions about how well or poorly the agent has comported himself *qua* agent. It also matters, independently, how other people were affected by the agent's faulty behavior, and how seriously they were wronged.

This brings a different solution to the puzzle of moral luck into view: Since an agent's level of blameworthiness is now understood as the degree to which she is an appropriate target for reactive attitudes like resentment and indignation, proposition (1) comes out as *true.* The two drunk drivers differ in their degree of blameworthiness, because of how their actions turned out.

But at the same time we can give an account of why (1) is true which shows that it is not, in fact, in tension with (2), the Control Principle. The Control Principle, recall, states that "an *agent* cannot be morally assessable for features of his action that were not under his control." But, crucially, blame and blameworthiness, as we are now understanding them, are *not* just a matter of 'agent-assessment'. In judging Ursula to be more blameworthy than Lara, we are not assessing Ursula as worse *in her capacity as agent*. What we are saying is that, because of the different ways in which their actions *affected* other people, they differ in their levels of blameworthiness – not because of, but *in spite of* facts about their relative levels of fault.

Understood in this way, there is a false presupposition in the puzzle of moral luck, as posed by Nagel: namely that when we blame one agent more than another, it is the *agent* who is differently assessed, depending on factors that are beyond her control.

A similar solution to the puzzle of moral outcome luck flows from a second prominent account of blame, namely T.M. Scanlon's. According to Scanlon, "[t]o blame a person for an action (…) is to take that action to indicate something about the person that impairs one's relationship with him or her, and to understand that relationship in a way that reflects this impairment."[73] Blame, on Scanlon's view, "is not a mere evaluation but a revised understanding of our relations with a person, given what he or she has done.

---

[72] Other prominent representatives of this Strawsonian tradition include R. Jay Wallace, *Responsibility and the Moral Sentiments* (Cambridge: Harvard, 1994); David Copp, "Defending the Principle of Alternate Possibilities: Blameworthiness and Moral Responsibility", *Noûs*, xxxi, 4 (December 1997): 441-56; and John Martin Fischer and Mark Ravizza, *Responsibility and Control: A Theory of Moral Responsibility* (New York: Cambridge, 1998).

[73] T.M. Scanlon, *Moral Dimensions: Permissibility, Meaning, Blame*, pp. 122-23.

Blame is therefore a function not only of the gravity of a person's faults but also of their significance for the agent's relations with the person who is doing the blaming." As Scanlon points out, how you, or a loved one, was affected by an agent's faulty behavior can multiply the *significance* of the fault for you, thus giving those impacted by her actions greater reason to revise their understanding of their relationship with her.[74]

Once again, therefore, the puzzle of moral outcome luck is solved by accepting (1), but in a way which doesn't challenge the Control Principle. The crucial difference between the lucky and the unlucky driver isn't how we should morally assess them *qua* agents. Rather, it is a patient-centric one: the way the unlucky driver's action has actually *affected* us give us greater reason to revise our understanding of our relationship with her.

10. Conclusion

The agent- and patient-centric evaluation of actions typically travel together. In most cases where an agent has comported herself poorly, this is connected to the fact that someone has suffered a moral injury at her hands; and the seriousness of the agent's fault will often correspond to how seriously others have been wronged. Likewise, if someone has been wronged, this will usually be at the hands of someone who has acted poorly.

If the arguments of this paper are sound, however, the two dimensions of evaluation sometimes part ways. Because we should be parochialists about the patient-centric notion of wronging, whereas we need not be parochochialists about the agent-centric notion of acting wrongly or being at fault, the agent- and patient-centric evaluations of action come apart in putative instances of moral outcome luck, like *Two Drunk Drivers*, and putative instances of moral dilemmas, like *Ticking Time Bomb*. This is what accounts for the distinctive phenomenology of these cases.[75]

---

[74] Ibid, especially pp. 149-50.

[75] These aren't the only cases where the two faces of morality come apart. As I have argued elsewhere, certain versions of Derek Parfit's Non-Identity Problem are cases where an agent acts morally wrongly but without wronging anyone. See my "Zukünftige Personen und Schuld ohne Opfer" ("Future Persons and Victimless Wrongdoing") in Markus Rüther and Sebastian Muders (eds.) *Worauf es ankommt: Derek Parfits praktische Philosophie in der Diskussion* (Hamburg: Felix Meiner Verlag, 2017).

What occurs in the cases of Lara and Ursula 'feels' different, because even though Ursula didn't comport herself worse *qua agent* than her lucky counterpart, seen from the patient-centric perspective, she *did a worse thing*. She grievously wronged a person, whereas Lara had a lucky escape.

		Agent-centric evaluation	
		The agent has comported herself well	The agent has comported herself poorly
Patient-centric evaluation	Other people have not been (seriously) wronged by the agent's conduct		Lara
	Other people have been seriously wronged by the agent's conduct		Ursula

Figure 1

As *Figure 1* illustrates, what makes no difference with regard to the agent-centric dimension of evaluation can make a big difference with regard to the patient-centric dimension. This is what accounts for our sense that Ursula has been unlucky.

Tragic moral conflicts arise because the converse is also true: As illustrated by *Figure 2*, what makes a crucial difference with regard to the agent-centric question whether an agent acted rightly or wrongly – namely whether his action was justified under a lesser evil logic – may not make a similarly decisive difference with regard to the patient-centric question whether the agent *wronged* those affected by his action.

		Agent-centric evaluation	
		The agent has comported herself well	The agent has comported herself poorly
Patient-centric evaluation	Other people have not been (seriously) wronged by the agent's conduct		
	Other people have been seriously wronged by the agent's conduct	**Large Man II; Ticking Time Bomb**	**Large Man I**

Figure 2

Thence the fact that although the agent in a case like *Ticking Time Bomb* or *Large Man II* may have acted flawlessly, indeed in the only permissible manner, he may nonetheless have caused another person a grave moral injury.

What emerges is an arrestingly split picture. Morality, it turns out, is Janus-faced: In our capacity as moral decision-makers, striving to act rightly and to avoid being at fault, we are in control of our destinies: I have argued that there is no need to accept that life and morality can conspire to place us in situations where we cannot but act wrongly. And even when we do comport ourselves poorly, by taking morally unjustifiable risks, how badly this impacts our records as moral *agents* is a function only of factors that were under our control.

Yet, in our capacity as inhabitants of a social world, who can relate to one another in various roles, including that of *wrongdoer* and *victim*, there is scope for great contingency and tragedy. We have seen that, when it comes to how seriously we have wronged another person by taking an unjustified risk, we are often utterly at the mercy of fate. And although there may be no genuine moral dilemmas, I have argued that life can nonetheless place us in situations from which we cannot escape without, in one way or another, feeling a profound sense of guilt.