

## The Power of Excuses\*

\* This paper has been presented at Birmingham, St Andrews, Oxford, Rutgers, and the 12th Conference of the Italian Society for Analytic Philosophy (SIFA). I thank the organizers and audiences for helpful discussions. I am particularly grateful to Bernhard Salow, Sarah Buss, Matthew Kramer, Christina Van Dyke, Sophie Horowitz, Katia Vavova, and Richard Holton who helped me improve previous drafts. Finally, I am very much indebted to two readers for this journals from whose generous and constructive comments this paper has greatly benefitted.

### 1. Introduction

Excuses are commonplace. Making and accepting excuses is part of our practice of holding each other morally responsible. But excuses are also curious. They have normative force. Whether someone has an excuse for something they have done matters for how we should respond to their action. An excuse can make it appropriate to forgo blame, to revise judgments of blameworthiness, to feel compassion and pity instead of anger and resentment.

The considerations we appeal to when making excuses are a motley bunch: tiredness, stress, a looming work deadline, a wailing infant, poverty, duress, ignorance. What unifies these various considerations as a class? In virtue of what can they all excuse? And what does their normative force consist in?

While particular excuses – chiefly, moral ignorance and duress – have attracted significant discussion in the literature, these more general questions have received less attention.<sup>1</sup> Indeed, some suggest that investigating particular excuses is all we can do; a unified account is not to be had. For example, Victor Tadros claims:

I do not think that it is particularly helpful to search for a general theory of excuses. Rather, we are best theoretically to investigate particular kinds of excuse claims. The upshot of such an investigation may be a general theory of excuses. However, I think that unlikely. I argue that there is probably no general theory of

---

<sup>1</sup> For discussions of moral ignorance, see Calhoun [1989], Moody-Adams [1994], Rosen [2004] and [2008], Harman [2015], [2011], Zimmerman [1988] and [1997], Arpaly & Schroeder [2014]. For discussions of duress, see Rosen [2014] and Baron [2014]. Discussion of excuses in general can be found in Austin [1957], Wallace [1994], Baron [2007], Kelly [2013].

excuses.<sup>2</sup>

I disagree. This paper aims to develop a unified account of excuses: what they are and what they do. In a nutshell, I argue that excuses are considerations that show that an agent's wrongdoing does not manifest a specific motivational failing: namely, the lack of a morally adequate present-directed intention. What do excuses do? I suggest that they function as responsibility-modifiers. They alter how the wrongdoer, the wronged party, bystanders may morally respond to a wrong, without negating that it remains appropriate to respond in some way.

Here is how the paper will proceed. I start by examining two influential accounts of excuses – one recently defended by Jay Wallace, the other tracing back to David Hume. I argue that, while initially promising, both accounts face significant problems. I then develop an alternative, the Good Intention Account, and showcase its advantages. A chief one, I argue, is that it captures the various considerations we appeal to when we make excuses into a unified account. Finally, I take up the relationship between excuses and moral responsibility. Contrary to received wisdom, I argue that excuses do not negate moral responsibility but modify the way in which an agent is morally responsible for her wrong.

A few clarifications before we start. To excuse can refer to an action: we excuse and ask to be excused. But we also say that someone has an excuse, that they are looking for an excuse, that something is (or isn't) an excuse. Used in this way, excuses refer to facts or considerations with particular normative significance. My focus here will be on this latter sense because I take it to be more basic. We ask to be excused by offering a consideration we take to constitute an excuse. We excuse someone when we think the consideration they put forward is an excuse.

The second clarification is this. We often talk dismissively about people “making excuses” when we want to say that they are offering “bad” excuses – spurious, unconvincing ones. Here, I will assume that a consideration is an excuse only if it in fact excuses. To make a “bad excuse” is to put forward a consideration that is not a genuine excuse.<sup>3</sup>

Third, a caveat. We make excuses for a broken promise but also for breaches of etiquette, for criminal offenses, for violations of professional norms, and even false beliefs. In this paper, I focus solely on excuses in our moral practice – whether we can extend the account to other domains will be left to future work.

---

<sup>2</sup> Tadros [2005], p. 295. Similarly, Baron frames excuses as having a disjunctive structure: they are either considerations that tell us something about the agent or about her situation. See Baron, 2007, p. 22-23.

<sup>3</sup> Relatedly, this paper is about the normative phenomenon of excuses, not our linguistic practice. Our everyday usage of the word “excuse” does not always line up with the normative phenomenon of excuses. To say “excuse me” can sometimes be a way of making a (passive-aggressive) accusation, rather than an excuse.

## 2. The Obligation Account

I will begin with an influential account of what excuses are and how they work. In general, we are to blame when our actions have violated a moral norm or obligation. According to a popular line of thought, excuses can render blame inappropriate in virtue of showing that although it may look like we have violated a particular moral norm, we haven't in fact done so. I will call this account the Obligation Account.

The Obligation Account has been most fully articulated and defended by Wallace.<sup>4</sup> Wallace argues:

...excuses function by showing that the agent did not really violate the moral obligations we accept after all. [...] To hold s morally responsible for x, when an excusing condition obtains, would involve the false belief that s's x-ing violated a moral obligation we accept; this gives us a reason for not holding people to blame when the excusing conditions are present.<sup>5</sup>

Why do excuses show that the moral obligation in question was not violated? Wallace accepts an account of moral obligation on which our obligations govern first and foremost our choices. Moral norms do not tell us what to do but rather what to choose. As Wallace argues:

...it is only through the mediation of our choices that the reasons expressed in moral principle may influence either our emotions or feelings, or the bodily movements we make. [...] one cannot be said to have *violated* a moral obligation in the absence of a relevant choice.<sup>6</sup>

Insofar as our moral obligations concern actions, they concern intentional actions: actions that were done as a result of a choice. Again, Wallace says:

More formally, one may say that moral obligations generally rule out doing actions of kind *x*, as a result of the choice to do something of kind *x*.<sup>7</sup>

The view that obligations govern our choices is not unique to Wallace. Thus, Barbara Herman argues that we should not conceive of moral obligations "as obligations to do certain actions (or kinds of actions)". Rather, she suggests:

The necessity that comes with an obligatory end constrains not

---

<sup>4</sup> Wallace [1994]. Hieronymi suggests a similar picture of excuses. Hieronymi argues that blame involves three judgments: (1) that the action is wrong, (2) that the wrongdoer is a member of the moral community, and (3) that you, the one who has been wronged, ought not to be wronged. Like Wallace, Hieronymi, too, regards a denial of (1) as an excuse. Unlike Wallace, she also classes the denials of (2) and (3) as excuses. See her [2001], p. 530.

<sup>5</sup> Wallace [1994], p. 133-34.

<sup>6</sup> *ibid*, p. 132, highlight in original.

<sup>7</sup> *ibid*, p. 144.

action but the will.<sup>8</sup>

The Obligation Account is simple and elegant. One of its virtues is a compelling explanation for why excuses shield an agent from blame. If blame responds to violations of moral norms and an excuse shows that the norm wasn't violated, then blame is undeserved. To continue to blame someone, while accepting their excuse, would be irrational.

Furthermore, it can seemingly accommodate the wide range of considerations to which we appeal when we make excuses. In particular, we can distinguish two types of excuses: first, considerations that show that the agent has simply not made any choice at all and thus cannot have violated a moral obligation. As Wallace argues:

...if one does x as the result of being pushed, or because of a muscular twitch or spasm, then one hasn't really acted at all.<sup>9</sup>

Second, there are considerations that concede that the agent made a choice but show that it was compatible with her moral obligations. Wallace continues:

...if one does x inadvertently, or by accident, then – though one may have acted – one didn't do x intentionally; and if one does x as a result of coercion or duress, then – though one may have done x intentionally – one hasn't merely done x intentionally, one has done x-rather-than-y.<sup>10</sup>

The Obligation Account accommodates why so many of our excuses appeal to an action's being unintentional: "I didn't mean to...", "I didn't realize...", "It wasn't my intention..." It says that there is a close connection between what we do intentionally and the choices our actions express. To dispute that an act was intentional is to dispute that the morally illegitimate choice was made and hence that a moral obligation was violated:

...the importance of intention lies in determining whether agent s has really done x, a morally impermissible act, in the first place. If the moral expectations we place on other people are primarily expectations concerning their attitudes toward us and others, as manifested in action, then what will be prohibited and required of people will not be types of bodily movement per se, but rather the attitudes expressed in bodily movements.<sup>11</sup>

The Obligation Account looks appealing. Its central thought is that excuses negate wrongdoing: their function is to show that an agent has not violated her moral obligations. This captures that many different considerations can comprise excuses, while identifying an underlying unity.

The Obligation Account faces two serious challenges. The first is to

---

<sup>8</sup> Herman [1993], p. 169.

<sup>9</sup> Wallace [1994], p. 121.

<sup>10</sup> *ibid*, p. 121.

<sup>11</sup> *ibid*, p. 126.

distinguish excuses from considerations that forestall the need for excuses. Wallace classifies considerations that show that an agent did not act, and hence did not violate her obligations, as a type of excuse.

But such considerations do not constitute excuses. Suppose that, trying to commit the perfect crime, you drug me unconscious, wrap my fingers around a gun, and pull the trigger, killing the victim. Since I was unconscious, I did not exercise my capacity for choice. I didn't act in any way. But it is odd to say that this gives me an excuse for the killing. Since I didn't act, I didn't kill. And so, I don't need an excuse. It's you, the murderer, who needs one.

The second challenge is to give a principled distinction between excuses and justifications. According to Wallace, considerations that show that the agent did act but her underlying choice was permissible and did not violate any moral obligations constitute a second type of excuses. Wallace suggests that this is how appeals to duress work:

[excuses appealing to duress work by] showing that agent's doing x actually expressed a different kind of motive: not merely a choice to do x, but a choice to do x-rather-than-y, or x-in-order-to-avoid-y. [...] Whether an explanation of this form will serve as an excuse will then depend on the content of our moral obligations-in-particular, whether they prohibit intentionally doing x-rather-than-y.<sup>12</sup>

When your choosing to do x-rather-than-y (or x-in-order-to-avoid-y) is permissible, you have an excuse.

This is not right. If, given your circumstances, it was genuinely permissible to make the choice you did, we do not invoke an excuse. Rather, we say that your action was *justified*. Suppose a villain puts a gun to my head, threatening to kill my child unless I compliment his sweater. The sweater is hideous, yet I oblige: I tell a lie. Given the circumstances, the lie is not just excused, it is justified. But if an action is justified, it needs no excuse.

Because the Obligation Account collapses the distinction between excuses and justifications it has trouble allowing that agents can have an excuse even when they act impermissibly. Consider someone who, subject to police intimidation, gives false testimony in court, which leads to the conviction of an innocent person. We can fill in the details to make it plausible that the action was impermissible: the interrogation is distressing but not physically violent, the threats made by the police are serious but not as serious as spending years behind bars. The agent ought not give in to the police intimidation; in doing so, she acts wrongly. Nevertheless, the coercive nature of the interrogation may constitute an excuse.<sup>13</sup>

More generally, tiredness, stress, and other emotional upset

---

<sup>12</sup> Wallace [1994], p. 144.

<sup>13</sup> See Rosen [2014] for another example of excused wrongdoing under duress.

plausibly do not bear on what moral obligations we have, nor do they prevent us from acting intentionally. Yet they can excuse. The Obligation Account cannot allow for that.

Views on which moral obligations concern choices, not actions, generally struggle to find conceptual room for a notion of excuse distinct from justification. A key motivation for such views is that it may not always be reasonable to expect us to perform some action but it is always reasonable to expect us to make the relevant choice. Thus, we can always be expected to abide by our obligations. This leaves no room for excuses, since those require the possibility of moral failure that does not reflect badly on us *qua* moral agents.

In what follows, I will assume that our moral obligations govern first and foremost actions, not choices. This is independently well-motivated. We can create new moral obligations by promising and waive existing ones by giving consent. An account of moral obligations needs to mesh with those moral practices.

Plausibly, the content of our moral obligations will mirror the content of what we promise and consent to. We consent to being operated on or hugged. We promise to pick up the child from nursery, to visit soon. We would be rightly suspicious if our partner phrased their promise in terms of choice rather than action: “Fair enough, you promise to choose to pick up the child. But do you also promise to pick up the child?” We promise and consent to actions, not choices.

### **3. The Character Account**

Excuses come into play when agents act wrongly: when they violate some moral norm or obligation. If excuses do not negate the charge of wrongdoing, they must negate some other moral criticism that can be inherent in blame. An influential alternative to the Obligation Account connects excuses to the wrongdoer’s character. It says that “an objectively wrong action (or an action in some way out of order) is excused if it *does not manifest some defect of character*.”<sup>14</sup> Call this the Character Account.

The Character Account improves on the Obligations Account in a few ways. It respects the distinction between excuses and justifications. Many considerations we offer as excuses – tiredness, severe stress – can be understood as circumstances in which the agent is “not properly herself”. But the Character Account faces trouble once we try to nail down its details.

A wrong action’s “not manifesting some defect of character” is not sufficient for its being excused. Character traits are dispositions that are made up of patterns of motivations across different times and circumstances. But when we ask whether an agent has an excuse for some wrong, we are not interested in how she *usually* acts. We are interested in

---

<sup>14</sup> Brandt [1969], p. 354. The inspiration for those accounts is Hume [1960], book II, part III, sec. II; and [1955], sec. VIII, part II.

her state of mind at the moment of action – how *it* was motivated. If you have betrayed my secret, it won't do to just point out that you are *generally* discreet.<sup>15</sup>

We can tweak the Character Account to escape the worry. Holly Smith suggests we understand it as follows: “An excuse functions, we might say, to block the natural inference from a wrongful act to the agent’s having a reprehensible motive for performing that act.”<sup>16</sup> Motive here refers to whichever psychological state motivates the action: a desire, intention, value, or affective state.<sup>17</sup> The focus on the motivations that leads to the wrongful action is an improvement. Still, the revised view faces difficulties.

Understood in this way, the lack of a reprehensible motive is not necessary for an excuse. Provocation can be an excuse. But when, incensed by my co-worker’s obnoxious remarks, I lose my temper and pour my cup of coffee into his lap, I am – at that moment – acting from a reprehensible motive. I desire to ruin his trousers and to scald his thighs. Losing my temper is a matter of this reprehensible desire “getting the better of me”. Still, the fact that I was provoked may constitute an excuse.

A lack of reprehensible motive is also not sufficient. Desires to kill, maim, harm, insult are clearly reprehensible motives. But much wrongdoing is caused by more mundane motivations: a student may cheat on her test to not disappoint her parents, a friend may supply a false alibi out of loyalty. Such motives may be morally criticizable but they are hardly morally reprehensible. Yet they do not constitute excuses.

The rest of this paper develops an alternative account of excuses: the Good Intention Account. I suggest that excuses point to the presence of a morally adequate motive, rather than the absence of a reprehensible one. Second, I suggest that it's a particular type of motive that matters for excuses: a morally adequate *present-directed intention*. The shift to present-directed intentions is crucial because present-directed intentions have a distinctive motivational profile that allows for a unified theory of excuses.

The next section sets out the Good Intention Account and its central ingredients.

#### **4. The Good Intention Account**

Here is the account, in a nutshell. When someone commits a wrong, we may generally infer that they lacked a morally adequate present-directed intention. Excuses block that inference. Excuses are

---

<sup>15</sup> See Gardner [2007], pp. 122-124 for further criticism.

<sup>16</sup> In Becker & Becker [2011], p. 507.

<sup>17</sup> *ibid.* See also Brandt [1969], p. 356 for this broad reading. Taking “motive” to include desires is established philosophical usage, see, for example, Arpaly [2003], Arpaly & Schroeder [2014], Markovits [2010]. Humeans, in particular, typically understand “motive” to denote all conative states partly because they take all conative states to be reducible to desires. See Smith [1994], pp. 92-125 for an articulation of this view.

considerations that show that the agent's wrongdoing does not reflect her lack of a morally adequate present-directed intention.

Let me start by unpacking the central ingredients. Following Bratman, Harman, Holton and others, I take intentions to be mental states characterized by their distinct functional role in action and planning.<sup>18</sup> One of their marks is their *stability*: once formed, they tend to persist, unless they are reconsidered and revised. A second mark is their close connection to action. They are *controlling*: unless revised, an intention will lead the agent to act.

Intentions can be future- or present-directed. Future-directed intentions commit us to plans of action in the future. A special class of such intentions are resolutions – these are standing commitments to courses of action. I may have a policy of only wearing black or not repeating gossip. Both future-directed intentions and resolutions facilitate inter- and intra-personal planning by curtailing deliberation. For example, if I have a policy of only wearing black, I do not need to expend cognitive resources on how to dress. I just reach for the black shirt.

Present-directed intentions concern what to do *right now*; they guide action in the moment.<sup>19</sup> The fact that present-directed intentions are controlling make them importantly different from desires.

I can have two simultaneous conflicting desires – to keep working on my paper and to go for a run – and act on one of them without revising or giving up the other. But I cannot have two simultaneous conflicting present-directed intentions and act on one of them without giving up or revising the other. Both are controlling and they both make inconsistent demands on my immediate actions. Unless I revise or give up one of them, I will experience practical paralysis: I cannot leave the house and remain seated at the desk at the same time.<sup>20</sup>

This has important consequences for which inferences we can draw about an agent's motivation, given her actions. Suppose you see me on a run outside. You are entitled to infer that I do not have a present-directed intention to keep working on my paper. For if I had *that* intention, I would

---

<sup>18</sup> See Bratman [1987], Holton [2009], Harman [1986] for arguments why intentions cannot be reduced to combinations of belief and desires.

<sup>19</sup> See Bratman [1987], Holton [2009], pp. 12-14 for discussion of present-directed intentions.

<sup>20</sup> Desires come in different strengths but intentions are widely seen as binary, so for two conflicting present-directed intentions there is nothing that could break the tie. The sole exception is Holton [2009], chapter 2 who argues for a notion of partial intentions. But Holton's sense of intentions being partial is very different from that of desires being strong or weak. Crucially, an agent with partial intentions has "two or more competing plans" *for the same end* (p. 36). Partial intentions make up a coherent plan with provisions for different ways the world might be. An agent forms partial, as opposed to all-out intentions in order to deal with epistemic uncertainty about which way of realizing the end will be effective.



be sitting at my desk. But you cannot infer that I have no desire to work on my paper. I may strongly desire to keep working but decided to take advantage of the brief sunny spell.

How do we acquire present-directed intentions? Future-directed intentions and resolutions will often automatically give rise to present-directed intentions. As Holton notes:

...if intentions are thought of as enduring states, then, for almost all intentions, there will come a point at which they are to be implemented. In that sense, then, they will be intentions to perform an action now. A future-directed intention will simply turn into a present-directed intention with the passage of time.<sup>21</sup>

Similarly, resolutions will turn into present-directed intentions when the relevant conditions obtain. A resolution to wear only black will turn into a present-directed intention to reach for the black shirt when I need to get dressed.

But we can also acquire present-directed intentions consciously by deliberating about what to do. When I think about whether to go to the Champion of the Thames or The Mill for an after-work pint and choose the latter, I thereby form the requisite present-directed intention to walk to The Mill. The ability to form present-directed intentions comes particularly handy in situations of incommensurability or indifference, where our desires are not enough to move us to act one way or the other.<sup>22</sup>

When it comes to excuses, it is present-directed intentions that matter. When confronted with wrongdoing, we care about the wrongdoer's motivation at the moment of action – not her general resolutions and plans. For this reason, and because the term "present-directed intention" is a bit of a mouthful, I will occasionally use the term "intention" and "present-directed intention" interchangeably. If at any point I mean to refer to future-directed intentions, I will flag this explicitly.

Suppose we witness a student cheating on her final exam to avoid disappointing her parents, thereby violating a moral obligation or norm. What are we entitled to infer about her motivation? We cannot infer that she was guided by a conative attitude with morally reprehensible content; it is not morally reprehensible to want to spare one's parents disappointment. Nor can we infer that she did not desire to be honest, for we often act in conflict with some of our desires. We can, however, infer that she is not guided by a morally adequate present-directed intention: she lacked the intention to be truthful. For if she had *that* intention, she would not have cheated.

In general, when we witness wrongdoing, we are entitled to infer the lack of a morally adequate present-directed intention. I say "in general" because this inference is defeasible. This is where excuses come in. Excuses are considerations that block this inference. An excuse shows that the

---

<sup>21</sup> Holton [2009], p. 13.

<sup>22</sup> *ibid*, p. 13.

agent's wrongdoing does not reflect the lack of a morally adequate intention. We have arrived at the Good Intention Account.

How do excuses block the inference from wrongdoing to the lack of a morally adequate present-directed intention? For limited agents like us, having a morally adequate present-directed intention is no guarantee that we will abide by our moral requirements. Several pitfalls remain. Our epistemic limitations may lead us astray; we have limited attention spans and time, incomplete (sometimes misleading) evidence, limited reasoning skills. We may prescribe a medical treatment to which the patient has an unforeseen adverse reaction or take a wrong turn on the way to the airport, causing our friend to miss her flight. Or we might stumble and lose our balance while helping our friend carry a valuable (and fragile) possession. Or a morally adequate intention can fail to get traction on our action because of limitations of our will: we have limited self-control, patience, tempers. In the face of terror, exhaustion, or provocation we may lose the ability to act on our intentions altogether.

Excuses thus work by offering an alternative diagnosis of our wrongdoing: rather than tracing it to a lack of morally adequate intention, they trace it to these other sources of error. This means that the Good Intention Account traces the structure of excuses to the structure of our agency and its limitations.

With the outline of the Good Intention Account on the table, the next step is to put it to work.

## 5. Unifying Excuses

This section considers the various considerations we put forward as excuses and argues that the Good Intention Account has a principled way of capturing them. At the same time, we will be filling in more detail.

*"I didn't do it on purpose!"*

Consider a common way of making an excuse: "I'm so sorry, I didn't mean to offend," or "I didn't intend to give away a secret – I thought this was common knowledge," or "I was just trying to...". They take the form of an appeal: yes, I did act wrongly – I broke the promise, gave away the secret, caused the offense – but *unintentionally* so.

The Good Intention Account explains why one's wrongdoing being unintentional can constitute an excuse. To say that it was unintentional is just to say that it did not reflect the lack of a morally adequate intention. Rather, something went wrong in putting my morally adequate intention into action: I slipped, tripped, lost my balance, or I was ignorant about some crucial feature of the situation.

Consider Jones who stirs a heaping spoonful of poison into Smith's tea. Since he has every reason to believe that the white powder is sugar, the poisoning is unintentional. In light of this, the wrongful deed does not reflect a lack of morally adequate present-directed intention. Jones is guided by the present-directed intention to sweeten Smith's tea; there is

nothing untoward about that. The poisoning reflects Jones' ignorance about the content of the sugar bowl.

"But wait," you might object, "surely, having done something unintentionally isn't always an excuse!" Indeed. The Good Intention Account does not say that *whenever* one's wrongdoing is unintentional, one is thereby excused. The considerations we appeal to in order to show that the wrongdoing was unintentional can sometimes themselves reflect a lack of morally adequate present-directed intention. Above, I stipulated that Jones' false belief about the white powder is justified. But suppose that it is the result of negligence – Jones pulled the package from the bathroom closet, not the kitchen, and didn't check the label, blithely assuming it to be sugar. In that case, Jones' false belief itself reflects a lack of morally adequate present-directed intention – namely, the intention to check the label carefully, so as to not put someone in harm's way. The Good Intention Account says, correctly, that it does not constitute an excuse. The same is true if Jones' false belief manifests recklessness. Thus, suppose Jones thinks there is some chance that the white powder is poison but refrains from investigating further. In this case, Jones deliberately took a gamble on Smith's life; that certainly reflects a lack of morally adequate intention. He has no excuse.

The Good Intention Account accommodates a common way of blocking the excuse that the wrongdoer was ignorant or misled about some morally significant feature of her action. We can block this excuse by pointing out that she *should have known*. This riposte does not deny that the wrongdoer was ignorant about the claim in question. But it maintains that the ignorance itself manifests a lack of morally adequate intentions.

### *Moral Ignorance*

Jones' wrongdoing resulted from him mistaking poison for sugar. But what if we are dealing with someone whose ignorance concerns a moral truth? What about a fervent Nazi who reports a Jewish family to the Gestapo, taking herself to "do her moral duty"?

Moral ignorance can be (and often is) the upshot of recklessness or negligence in one's moral inquiry; as such, it fails to excuse. Still, there are environments in which certain kinds of wrongdoing are so normalized that it is plausible an agent may have acquired a mistaken moral belief non-culpably.<sup>23</sup> The question for the Good Intention Account is whether someone can both act on a false moral belief and yet have a morally adequate intention.

It may be tempting to think that false moral beliefs about a situation will always prevent an agent from forming a morally adequate intention in that situation. The fervent Nazi's misguided beliefs in Nazi ideology prevents her from forming the intention to respect all persons' humanity. Her wrongdoing reflects this lack. This suggests that wrongdoing done

---

<sup>23</sup> See Calhoun [1989].

from moral ignorance cannot be excused.

But this is too quick. Sometimes agents act on false moral beliefs and we do regard their wrongdoing as excused. Consider an overworked emergency doctor who genuinely cares about her patients. She is confronted with a stroke patient and has to decide between two treatment options, trading off risk of death against expected quality of life in a snap decision. She decides to pursue the more invasive treatment because she gives slightly greater weight to minimizing risk of death. Suppose that's a mistake. She ought to have given greater weight to quality of life and pursued the less invasive treatment – her moral deliberation was flawed.

The doctor's wrong action is based on a mistaken moral belief. Still, the fact that she tried to do what's best for the patient counts in favor of cutting her some slack. We see her wrong action in a different moral light than that of the fanatical Nazi or that of a cynical doctor who pursues the more invasive treatment because it allows her to charge the patient more.

The challenge is to reconcile these seemingly conflicting intuitive verdicts: the Nazi's belief that she's "doing her moral duty" does not move us. Yet the fact that the doctor is trying to do what's best for the patient does.

I think reconciliation is possible. To see how, we need to say more about moral adequacy and the structure of present-directed intentions.

First, intentions. Present-directed intentions are plans that guide actions. This means that their content cannot be exhausted by an abstract goal. I cannot act on a present-directed intention whose sole content is "to go to Japan" because this is not specific enough to determine a course of action. My intention needs to include a conception of how to realize that goal. For example, I intend to go to Japan *by* taking a flight out of Heathrow *by* boarding the 16:30 British Airways flight to Tokyo. The content of intentions has hierarchical structure: there's an overarching goal and a conception of the steps necessary to realize it.<sup>24</sup>

This suggests a natural way of how to evaluate intentions for adequacy. To assess an intention for adequacy is to assess it as a plan with respect to a given goal. How adequate a plan is with respect to that goal depends on how conducive it is to realizing it – for example, how conducive my intention to go to Japan is to getting me there.

The aim with respect to which we evaluate an intention or plan will often be internal to it. But we can also evaluate an intention with respect to an external goal. We might evaluate my intention to go to Japan with respect to how conducive it will be to my health or to minimizing my carbon footprint, even if those goals does not figure in my intention at all.

Adequacy is not a binary notion: given some particular aim, plans can be more or less adequate. A plan can be, on the whole, adequate even when it is ultimately unsuccessful. My plan to go to Japan by taking the

---

<sup>24</sup> This is congruent with work on action-guidance by neuroscientists. See, e.g. Duncan et al. [1996].

train to Heathrow may be adequate even when it fails – there’s an accident on the line and train service to London Heathrow is suspended. Contrast this with the plan to go to Japan by taking the Eurostar (based on a poor grasp of geography) – that plan is clearly inadequate; it is doomed from the outset.

Now for *moral* adequacy. We can evaluate intentions for moral adequacy by assessing them relative to a relevant moral goal – for example, the goal of doing the right thing, or treating all persons with humanity. With this in mind, let us revisit our cases above. Even if the Nazi’s overarching aim was to “do the right thing”, this could not have been the sole content of the intention guiding her. “Doing the right thing” is too abstract to issue in action in any given situation. The full content of the intention guiding the Nazi is plausibly something like “to do the right thing *by* ensuring the racial purity of her community *by* weeding out all Jews.” And even if the Nazi’s own overarching goal of “doing the right thing” is morally fine, the full intention is morally objectionable.<sup>25</sup> Keeping one’s community “racially pure” is incompatible with doing the right thing, not just in this particular situation but in all conceivable circumstances – the Nazi is spectacularly mistaken about the nature of doing the right thing. And, of course, the Nazi’s present-directed intention also fails to be morally adequate when evaluated with respect to other moral goals: treating all persons with humanity, for example.<sup>26</sup>

Contrast this with the doctor. She acted on the present-directed intention to do what is best for the patient *by* pursuing the treatment option that fares best on a careful analysis of risks and benefits *by* pursuing the one that fares better on survival rates. Doing what’s best for the patient is a morally sound aim and carefully trading off risks and benefits of the treatment options is a good way to realize that aim. The doctor is correct about which facts about the treatments are morally relevant and their moral valence. Where she goes wrong is in the very final step: in exactly how to weigh them.

In light of this, we may well judge the doctor’s present-directed

---

<sup>25</sup> Those moved by “moral fetishism” worries will think that there is something morally objectionable about such an overarching goal or any intention with “de dicto” content. See Smith [1994], p.98, Arpaly [2003], pp. 71-73, Markovits [2010], p. 203. They see other moral aims as relevant, such as respecting all persons’ humanity.

<sup>26</sup> Could the Nazi have two conflicting present-directed intentions, for example, to “eliminate all Jews” and “treat all persons with humanity” but only act on the former? No, for recall that present-directed intentions are controlling. Unlike desires, present-directed intentions do not come in different strengths, so there is nothing to break the tie in the face of such a conflict. An agent with two conflicting present-directed intentions experiences practical paralysis until she revises or gives up on one of them. It is, of course, possible that the Nazi acts on the present-directed intention to “eliminate all Jews” while also having a *future-directed* intention to treat all persons with humanity. Such a combination of attitudes is psychologically possible, albeit arguably irrational. But future-directed intentions are irrelevant to excuses. So the future-present conflicted Nazi does not have an excuse. I am grateful to an anonymous referee for raising this objection.

intention as, on the whole, morally adequate, particularly in the context of a stressful overnight shift, which suggests that the doctor was being neither negligent nor reckless in her deliberation. To say that the doctor's present-directed intention is, on the whole, morally adequate is not to say that it is *wholly* morally adequate. After all, it contains a moral mistake. But it contrasts with the cynical doctor who pursues the more intensive treatment to maximize the patient's hospital bill. The cynical doctor's actions reflect a lack of a morally adequate intention – maximizing the patient's hospital bill is not a way of doing what's best for the patient. Nor is it morally adequate when evaluated with respect to other plausible moral goals: doing the right thing or respecting all persons' humanity.<sup>27</sup>

I suggested that we evaluate an agent's intention for moral adequacy by assessing how adequate it is as a plan with respect to a relevant moral goal. This leaves a crucial question unaddressed, you might say: namely, what the criterion for moral adequacy is. What is the relevant moral goal? And how adequate is adequate enough to ground an excuse?<sup>28</sup>

The Good Intention Account does not aim to settle these questions. Different normative theories will disagree about what the relevant moral goal is – whether it is to maximize pleasure and minimize pain in the world or respect the humanity of all persons. They may disagree about whether there is a unique moral goal. It's not for the Good Intention Account to settle those disputes. The Good Intention Account predicts that occasionally different normative theories can come to different verdicts on whether some consideration constitutes an excuse.

Equally, there can be moral disagreements about whether an agent's intention is sufficiently adequate to furnish them with an excuse. Again, these will have to be settled through first-order normative inquiry; I doubt that we can write down a simple formula.

The Good Intention Account diagnoses why disagreements about whether some wrongdoer has an excuse can persist, even when both parties agree on all features of the case, including the wrongdoer's motivations. Such disagreements reflect conflicting moral judgments about the moral adequacy of an agent's present-directed intention. I complain that your family member is constantly meddling, you say they are merely trying to help. I may well agree that their present-directed intention is to help but

---

<sup>27</sup> The thought that moral adequacy admits of degrees chimes with the observation that excuses can come in different strengths: given a particular wrong, some excuses are better than others. One factor that can determine the strength of an excuse is *how* morally adequate it shows the agent's present-directed intention to be. This may also explain why, other things being equal, moral ignorance generally strikes us as a weaker excuse than nonmoral ignorance. A present-directed intention that includes a moral misconception will necessarily fall short of full moral adequacy. These suggestions require much more careful spelling out – something that I will have to leave for another paper. See my [forthcoming]. Still, the fact that the Good Intention Account has resources to make sense of excuses coming in different strengths is reassuring.

<sup>28</sup> I'm grateful to an anonymous referee for pressing me on this.

think that that's not a morally adequate intention to have. The morally adequate intention, in my view, would be to stay out of it – a present-directed intention the family member decidedly lacks.

The Good Intention Account thus says that a morally ignorant agent can sometimes have an excuse without saying that all instances of wrongdoing done from moral ignorance are excused.

### *Moral Dilemmas*

The Good Intention Account explains why an agent's being caught in a moral dilemma constitutes an excuse. Such an agent is forced to choose between two morally impermissible choices: she will violate a moral norm or obligation whatever she decides on. On the Good Intention Account, it's not surprising that we regard such an agent as having an excuse. If the situation is one in which she will act wrongly no matter what she does, her wrongdoing reflects her tragic options, not the lack of a morally adequate intention.<sup>29</sup>

### *Emotional Upset*

Another important class of excuses includes duress, provocation, and other emotional upset. The agent who gives in to police interrogation and falsely testifies against an innocent person acts wrongly under duress. The agent who snaps at a colleague because she is dealing with difficulties at home acts wrongly because she is stressed. Duress, coercion, stress do not generally render the agent's actions unintentional. Nor do they need to compromise the agent's epistemic capacities: the agent may give in to the interrogation, knowing that she shouldn't or snap at her colleague, knowing it is an unkind thing to do. Yet they can excuse.

I suggest that when a stressful situation, emotional upset, or tiredness excuse, it is because they cause the agent to lose self-control. Self-control is the ability to put one's present-directed intentions into action in the face of contrary desires. With her self-control lost, an agent cannot hold these contrary desires in check. Hence, she directs a sharp comment at her colleague, or gives in to the police's demand.<sup>30</sup>

Loss of self-control excuses because it prevents the agent's intention from guiding her action. The agent who loses self-control does not lose or revise her intention. It's a mark of loss of self-control that one's intentions remain intact – what goes missing is one's ability to act on them. Instead, with her self-control undermined, the agent under stress or duress is guided by her momentary desires and inclinations: the desire to shut up the colleague, or to just make the interrogation stop.

Let me say a bit more about what is involved in losing self-control

---

<sup>29</sup> Of course, provided she didn't wind up in the dilemma through negligence or recklessness.

<sup>30</sup> See Holton [2009], chapter 6 for a detailed discussion of strength of will.

and how it differs from weakness of will.<sup>31</sup> To illustrate, consider the well-known discussion by Austin:

I am very partial to ice cream, and a bombe is served divided into segments corresponding one to one with the persons at High Table: I am tempted to help myself to two segments and do so, thus succumbing to temptation and even conceivably (but why necessarily?) going against my principles. But do I lose control of myself? Do I raven, do I snatch the morsels from the dish and wolf them down, impervious to the consternation of my colleagues? Not a bit of it. We often succumb to temptation with calm and even with finesse.<sup>32</sup>

The phenomenology of losing control of oneself – the experience of snapping at a colleague, while overtired, or giving away a secret when put under intense pressure – is very different from that of giving in to the temptation of a second piece of cake. This difference is reflected in how we talk about the two phenomena. We “give in” to temptation, we “are (or let ourselves be) seduced”, we “succumb” to weakness of will. In contrast, to lose self-control is “to lose it”, “to break down”, “to explode”, “to snap”.

The difference in phenomenology is underpinned by a difference in psychological mechanism. As Holton writes:

...a person exhibits weakness of will when they revise an intention in circumstances in which they should not have revised it.<sup>33</sup>

Thus, Austin’s desire for ice cream leads him to reconsider and give up on his intention to only take one piece, forming the intention to help himself to two slices instead. This new intention guides his action. The weak-willed agent’s ability to form and be guided by her intentions remains intact. Some ways of giving in to temptation can require a great deal of planning and, in Austin’s word, “finesse” (waiting for an opportune moment to swipe the second slice from the platter). The revising of our intention is often preceded by deliberation: we reconsider our options, we make up justifications and excuses. (Shouldn’t Professor Jones abstain, given her high cholesterol? Hasn’t the Bursar stealthily helped himself to two pieces of cake the other day?)

In contrast, the agent who “snaps” does not revise her intention to stay calm and polite to her colleague. Rather, her intention has simply no traction on what she does – she is so upset that her ability to plan is bypassed altogether, with her momentary desires winning out. Such loss of ability affects the agent in more than just a “localized” way: not only does the intention to stay polite have no grip, neither does the intention to stay out of trouble, to not be found out, etc. As Holton & Schute argue:

[W]e would expect a loss of self-control to have a systematic effect

---

<sup>31</sup> In this I am following the view set out by Holton [2009].

<sup>32</sup> Austin [1957], p. 24.

<sup>33</sup> Holton [2009], p. 73.



on those actions that it would ordinarily inhibit. So, for instance, a lack of concern on the part of the agent for her own long-term welfare (manifested, say, in a lack of concern about being caught) would be evidence that self-control was in fact lost; in contrast, evidence of such concern (donning gloves before picking up the hammer to ensure that no finger prints were left) would be evidence that it was not.<sup>34</sup>

The weak-willed agent revises her intention when she shouldn't, retaining self-control. The agent who "loses it" retains her intention but loses her ability to act on it.

Self-control is not the same as the ability to act intentionally. Thus, to say that Sara poured coffee into her colleague's lap because she lost self-control is not to say she did so unintentionally. The coffee pour was not just an unlucky accident – it was motivated by her momentary desires and affective states: her anger at his obnoxious comment and her desire to shut him up. Nevertheless, her pouring the coffee is compatible with her having a present-directed intention to keep her cool. The provocation and the resulting anger made it impossible for that intention to guide her action.<sup>35</sup>

You might worry about the claim that when losing self-control we are "swept away" by our momentary desires. Are agents who lose self-control really unable to reign themselves in? Take an exhausted parent, who shouts at her toddler. She may well have "kept it together", had she tried harder, for example, had she been out in public. This may suggest that it's never impossible to retain composure – it just becomes harder.<sup>36</sup>

In response, we need to distinguish the loss of self-control from the process leading up to it. Before self-control is lost, it is generally eroded. The more your self-control is eroded, the more effort you have to put in to not lose it and the wider the range of situations and circumstances in

---

<sup>34</sup> Holton & Schute [2007], p. 61.

<sup>35</sup> I am assuming here that what Bratman [1987] calls The Simple View – the claim that "A  $\varphi$ -ed intentionally if and only if A acted on the intention to  $\varphi$ " – is false. This assumption has been subject to extensive discussion in the philosophy of action literature. See also Bratman [1987] chapter 8. See Holton [2009], p. 15-17 why The Simple View should be rejected.

<sup>36</sup> I'm grateful to an anonymous referee for raising this objection.

which you are vulnerable to losing it.<sup>37</sup> For example, the greater the agent's fear, stress, or sleep deprivation, the more effort she has to put in to resist the police's aggressive questioning and the more likely she is to "crack" as additional pressure is applied. The more sleep-deprived the parent, the more difficult it is for her to remain patient at her toddler's meltdowns and the more external support she requires to "keep it together".

The objection is right that in many cases in which the agent loses self-control, she may have kept her composure, had she put in more effort. But how much effort we are capable of depends on a variety of factors, including whether one is well-rested or out in public. To say that one could have tried harder in a counterfactual situation is not to say that one could have tried harder in the actual one. Given the parent's actual circumstances – sleep-deprived, home alone – greater effort may have been beyond her.

Even if one could have avoided losing self-control with greater effort, *once* self-control is lost, one is at the mercy of one's desires. It's only at that point – once we "lose it" – that all effort gives out and we are overcome by an "ungovernable passion". Loss of self-control is all-or-nothing and instantaneous – from one moment to the next, the agent "snaps", "breaks down", "explodes". But, as already noted, the lead up to this point can be a drawn out, gradual affair – just as a balloon stretches further and further until, eventually, it explodes.

It's important that gradual erosion of self-control often precedes its loss because it means that loss of self-control can be culpable. Like mittens and diamond earrings, tempers can be lost through negligence or recklessness. We can knowingly put ourselves in situations that will erode our self-control and knowingly fail to extricate ourselves from situations that will erode it further.

We can be complicit in the erosion of our self-control by picking fights, drinking too much alcohol, neglecting to change topics of conversation, skipping lunch, failing to step out of the room for a deep breath or two. But our self-control can also be eroded through no fault of our own: a teething baby who is up all night, a family disaster, an unmanageable workload.

Let us return to wrongdoing. One virtue of the Good Intention Account is that it vindicates that weakness of will is no excuse. While both loss of self-control and weakness of will can lead to wrongdoing, they do

---

<sup>37</sup> These observations are congruent with experimental work in neuroscience and psychology. There is currently an ongoing and lively empirical debate about the mechanism of self-control. Notable contenders include Inzlicht's attention-based model (see Milyavskaya & Inzlicht [2018]) and Baumeister's resource-based model (see Baumeister [2000]). Both bear out the folk-psychological thought that self-control can be eroded but have different implications for how erosion works, including whether depletion effects are domain-general (i.e. whether an agent who has exerted self-control on one task will show depleted self-control on an unrelated task, as defended by Baumeister) or task-specific (as favored by Inzlicht). See Friese et al. [2019] for an overview over recent empirical evidence, which argues that empirical results are inconclusive.

so in different ways. The weak-willed agent acts wrongly because she lacks a morally adequate intention – her weakness of will consists in her having given up on it. And so, according to the Good Intention Account, she does not have an excuse.

In contrast, the agent who as a result of duress, coercion, tiredness, or other distress loses self-control acts wrongly because she momentarily loses her ability to translate her present-directed intentions into action. Wrongdoing that results from loss of self-control thus need not reflect the lack of morally adequate intention.

The Good Intention Account says that loss of self-control can excuse, not that it always does. The restrictions on loss of self-control as an excuse mirror those on ignorance. Neither excuses when it's the upshot of negligence or recklessness. This might explain why we typically look more favorably on the excuse of duress or coercion than on provocation. An agent under duress typically has not had any say over her upsetting situation. In cases of provocation things are often less clear cut. Even if we trust that the agent genuinely lost self control in the course of the bar fight, we might question whether she is culpable for having found herself in the midst of it.

Let's step back. I argued that the Good Intention Account captures the motley bunch of considerations we recognize as excuses into a unified account: it traces the structure of excuses to the structure of our agency. Excuses which appeal to one's wrongdoing being unintentional, trace the wrongdoing to limitations that are epistemic or agential (clumsiness, loss of balance). Excuses that appeal to one's wrongdoing being the result of duress, coercion, tiredness, stress, etc, trace the wrongdoing to limitations in our will. Both make the case that, given the circumstances at hand, the wrongdoing does not indicate the lack of a morally adequate present-directed intention.

The next task is to see how the Good Intention Account responds to some objections.

## 6. Making Excuses and the Good Intention Account

The first worry is that the Good Intention Account makes excuses too individualistic. It says that whether someone has an excuse is a matter of *their* moral psychology: the presence or absence of morally adequate intentions. But, you might object, when we consider whether a wrongdoer has an excuse, we draw on general expectations about agents' range of capacities. Our practice of excuses is tied to this range.<sup>38</sup>

I suggest that general expectations about agents' range of capacities play a crucial *epistemic* role in our practice of excuses. The reason why we dismiss out of hand considerations that completely contravene these general expectations – that Jill punched Samantha because she missed 10 minutes of sleep last night – is not that such considerations could not

---

<sup>38</sup> I'm grateful to an anonymous reviewer for raising this objection.

possibly excuse but rather that they suggest a moral psychology that is so unusual as to be beyond belief. (Not even toddlers are that fussy about 10 minutes of sleep.)

In support of this, note that our assessment of excuses becomes both more nuanced and idiosyncratic the more familiar we are with someone and their quirks. We might not take offense at our partner's snide remark in the car because we know that she gets particularly stressed out driving in heavy traffic. We recognize this as an excuse for her misstep even though we would not recognize it as an excuse in general. Equally, knowing someone very well, we might not accept a consideration as an excuse, although we would do so for a stranger. If I know that Bertha has nerves of steel and thrives on the adrenaline of 5 simultaneous deadlines, I will not regard overwork as a plausible excuse for some misdeed of hers, even if I would generally cut someone in her position some slack.

When we offer a highly idiosyncratic excuse, we do not just face the challenge of making it believable but also of establishing that, e.g. the loss of self-control was non-culpable. Suppose I can make a credible case that I snapped at my colleague because of caffeine withdrawal to which I am extremely sensitive. But if I know that missing one cup of coffee makes me so vulnerable to "losing it" – as I must, if I offer it as an excuse – why didn't I take precautions? Why didn't I hold off that conversation until I caffeinated myself? Why don't I keep an emergency jar of instant coffee in my house?

A second observation is that our practice of making excuses often appeals to the relevant offense being "out of character" for the wrongdoer. Can the Good Intention Account explain this? On the Good Intention Account, such an appeal does not function as an excuse but rather as an expression of confidence that *there must be* some excuse. Since we know this person always to be guided by morally adequate intentions, the only way we can make sense of her wrongdoing is that something else must have gone awry. She must not have realized what she was doing, or must have lost self-control.<sup>39</sup>

Third, you may worry that the Moral Intention Account is vulnerable to counterexamples. In particular, you may think that duress and coercion can excuse even when there is no loss of self-control. Suppose someone threatens to harm my child unless I steal your wallet. I coolly calculate the risks involved and decide to commit the theft. Doesn't the presence of the threat constitute an excuse even when my self-control was entirely unaffected? After all, I can bring up the threat in my apology and

---

<sup>39</sup> The Good Intention Account also captures that appeals to character can block excuses. You say I need to cut Mike some slack for being mean – he was stressed about a deadline. I point out he's always a jerk, not just when he's stressed. Thus, there is good independent reason for thinking that the agent generally lacks the relevant morally adequate intention and thus no reason to think it was present (but short-circuited) in this case.

expect it to mitigate my guilt.<sup>40</sup>

According to the Good Intention Account, the mere presence of a threat does not constitute an excuse – what matters is its effect on your motivation. To furnish you with an excuse, it must have either affected your ability to reason or undermined your self-control (or both). If it did neither, it does not constitute an excuse. Still, the threat could make a difference to whether you are to blame because it's plausible that, given the stakes, the theft was *justified*: the threat rendered stealing your wallet permissible.<sup>41</sup>

Excuses and justifications both bear on whether (and how much) an agent is to blame but they do so in different ways. Justifications bear directly on the moral status of the action: on whether it was in fact wrong or impermissible. Excuses, in contrast, bear on the agent's motivation. To make an excuse is not to deny that one's action was wrong. Rather, it's to point to one's presence of morally adequate intention in order to mitigate how much one is to blame.

The Good Intention Account is an account of excuses, not of justifications. But some have suggested that the notion of excuse is derivative of the notion of justification. If so, it's a mistake to think we can talk about one without talking about the other. Gardner has defended such a view; he argues that justifications are reasons why a given action is permissible; excuses are justified but false beliefs about those reasons:

[Excuses] are what the person acting upon them takes to be valid reasons, *and justifiably so*. Thus the structure of excuse derives from the structure of justification. To excuse an action is not, of course, to justify that action. Rather, one justifies one's belief that the action is justified.<sup>42</sup>

But there are reasons to doubt that the connection between excuses and justifications is so tight.

First, some considerations excuse in virtue of interfering with the agent's reasoning, putting a justified belief beyond her reach. Suppose you put me under enormous stress as I am making a decision with important moral implications. I mess up my deliberation and, as a result, do the wrong thing. In appealing to this severe stress, I need not claim that I came to a justified belief that the course of action I took was permissible. I can simply point out that, given your interference, I was not in a position to come to a justified belief about what to do.

---

<sup>40</sup> I'm grateful to an anonymous referee for raising this objection.

<sup>41</sup> The Good Intention Account also accommodates excuses for failures to perform supererogatory actions. In general, we can infer from someone's failure to perform a morally good (but not morally required) action that she lacked the present-directed intention to do what's good. An excuse for failing to do the morally good (but not morally required) thing shows that this inference is unwarranted. For example, I didn't give the homeless man change because I didn't have my wallet on me.

<sup>42</sup> Gardner [2007], p. 110, my highlight.

Second, Gardner's view does not sit well with many excuses that involve loss of self-control. Losing self-control can be understandable, relatable or regrettable but it is not something that can be reasonable or justified. Still, Gardner suggests that loss of self-control can be "reasonable" insofar as it is lost as a result of beliefs or emotions that are "reasonable". We make an excuse for losing our temper by justifying the emotion or belief that led to it.

By 'reasonable' here is meant, in my view, much the same as 'justified'. There must have been an *undefeated reason* for one's belief, emotion, etc. which also explains why one held the belief or experienced the emotion, etc.<sup>43</sup>

This may be plausible for anger or fear but not for tiredness, nausea, hunger, which we also recognize as excuses. Insofar as we think that the parent's sleep-deprivation constitutes an excuse for her yelling, it's not because we take her sleep-deprivation to be justified; tiredness, like nausea and headaches, does not admit of rational evaluation. Nor do we have to attribute to her a justified belief that it was permissible to yell. When these considerations excuse – that is, when the agent is not culpable for finding themselves in the relevant state – it's because they lead the agent to lose self-control.

The structure of excuses does not mirror the structure of justifications. Rather, it mirrors the fault lines of our agency. Excuses show that an instance of wrongdoing does not reflect a particular failing of our moral motivation: a lack of morally adequate intention. In doing so, the excuse need not portray the action as justified or reasonable by the agent's own lights. As Holton & Schute argue for loss of self-control:

...while the loss of self-control is not reasonable, it is something that can happen to the reasonable person.<sup>1</sup>

Their point generalizes: mis-evaluating one's evidence or losing one's balance is not reasonable but these things, too, can happen to a reasonable person. Reasonable humans are still human; they remain susceptible to making mistakes when tired, stressed, anxious, to fall for misleading evidence, to lose their temper when angry and their nerve when afraid.

The Good Intention Account gives us an understanding of what excuses are – in virtue of what considerations excuse. The question I turn to now is what excuses *do*.

## 6. Excuses as Responsibility-Modifiers

The received wisdom says that an excused action is one for which the agent is neither to blame nor responsible. Thus, take Wallace:

Excuses [...] aim precisely to challenge the claim (or suspicion) that S was morally responsible for x.<sup>44</sup>

---

<sup>43</sup> *ibid*, p. 112.

<sup>44</sup> Wallace [1994], p.121 The same line is also expressed by Murphy [1982], p. 508, Allais [2008], p. 34-35. Rosen [2008], p. 609, Baron [2007], p. 27, Hieronymi [2001], p. 530.

Excuses, on this view, are powerful indeed; they demarcate the boundaries of moral responsibility.

I disagree; I argue that excuses do not negate but modify an agent's moral responsibility. They bear on the way in which an agent is responsible for what she has done.

Before I present the positive picture, I have to dislodge the received view. My argument that excuses leave moral responsibility intact goes as follows: excuses leave moral residue; they leave something for the agent to apologize, to make amends for, to feel distress about. The best explanation for why excuses leave such moral residue is that the agent is morally responsible for the wrong done.

Moral residue comes in two kinds: psychological and normative. Psychological residue refers to the attitudes and feelings rendered apt by wrongdoing: these can include feelings of guilt, remorse, anguish, sorrow, and grief on the part of the wrongdoer and resentment, indignation, disappointment on the part of victims and onlookers. Normative residue refers to the duties that one acquires as a result of having acted wrongly: to acknowledge what one has done, to apologize, to make amends, to explain oneself, to compensate. These duties – reparative duties, as we may call them – are duties the wrongdoer owes to the wronged party, who holds a correlate set of reparative rights against the wrongdoer.

Lack of moral responsibility entails lack of moral residue. To see this, consider an example: your house mate accuses you of breaking a tea mug she held particularly dear and attempting to hide this from her (she found the shards in the trash). Your conscience is clear: someone else was the culprit. In light of this, you are not morally responsible. Consequently, it would be inappropriate for you to feel guilt, remorse, or sorrow. You may regret the sorry state of affairs: the mug's tragic demise, your housemate's upset. But you do not have to regret anything that *you* have *done*. You may owe it to your friend to cheer her up – but that's a duty of friendship, not a reparative duty. You do not owe your friend an apology, a new mug, nor some other “making it up”-gesture. There is neither psychological nor normative residue.

Things are different when it was *you* who broke it – even when you have an excuse. You were careful but the mug slipped out of your soapy hand while doing the dishes. You didn't know that your friend was particularly attached to it (it was ordinary and in communal use in the kitchen), so you simply cleaned up the mess. Still there is both psychological and normative residue in this case: it is appropriate for you to feel bad for having broken the mug, you owe your friend an admission of guilt, an apology, an explanation for why you didn't tell her, perhaps even a replacement mug. If you refused to apologize, simply shrugging and insisting that this was an accident and you are faultless, your friend would rightly be upset. By failing to apologize, you would violate a reparative duty and thus give her grounds for further complaint, over and above the original one.

If excuses negated morally responsibility, we would expect the two cases to be morally equivalent. But they are not: there is moral residue in one but not the other. What explains this asymmetry? A natural explanation is that in one case you are morally responsible for breaking the mug – although you have an excuse. In the other you are not.

Let me fill in the resulting picture of moral responsibility a bit more before addressing objections.

Following Strawson, a common approach to moral responsibility views it as closely connected to a narrow set of reactive emotions: anger, indignation, and resentment. Thus Wallace argues:

When we hold a person responsible for an unworthy act, we are subject to a negative reactive emotion because we believe the person to have violated a moral obligation we accept, or we believe that such an emotion would be rendered appropriate by the violation.<sup>45</sup>

On this view, the paradigm way of holding someone responsible for a wrong is to resent them for it (or to regard such resentment as apt in light of the wrong).

My suggestion is to broaden our view of moral responsibility to include not just emotional responses and their expressions but also normative residue. Resentment is just one kind of psychological residue left by wrongdoing – others can include disappointment, sadness, disillusion, frustration, pity. And wrongdoing leaves normative residue, too: owing an apology, an explanation, compensation.<sup>46</sup>

This suggests that moral responsibility can manifest in many different ways.<sup>47</sup> Such a view is independently attractive; it fits naturally with how we think of moral responsibility for right actions. Strawson highlights gratitude as a reactive attitude that responds to right actions.<sup>48</sup> Being grateful for someone's kind action is certainly a way of giving them credit for it. But gratitude does not have a monopoly here: there is approval, admiration, respect, awe. There is owing someone some kind of acknowledgement or a favor in return. I can hold you morally responsible for your good deed by regarding any one of them as warranted.

---

<sup>45</sup> Wallace [1994], p. 71. Wallace then goes on to narrow down the relevant class of negative reactive emotions to resentment on p. 76.

<sup>46</sup> Note that while these attitudes are typically rendered apt by wrongdoing, some of them can be apt even when there was no wrongdoing because the agent's action was justified. It may be inappropriate to feel resentment and outrage towards such an agent, still other responses remain apt: she may owe an explanation and apology, it may be apt for her to feel bad. This fits well with the observation that justifications do not negate moral responsibility – when we justifiably violate a moral obligation, our action leaves a moral remainder.

<sup>47</sup> Smith [2015], p. 103 suggests a similar picture, though her way of spelling it out differs from mine.

<sup>48</sup> Strawson [1962].



On such a view, excuses function as *responsibility-modifiers*. They change the psychological and normative residue that remains in the wake of a wrong; they modify which attitudes are appropriate and which reparative actions are required. Excuses bear on the specific way in which an agent is morally responsible for what she has done; they bear on how wrongdoer, wronged party, and the moral community may and/or should respond to it. The normative effect of any excuse will depend on the details of the case: the wrong in question, its circumstances, the nature of the relationship between the two parties. But excuses do not negate moral responsibility: some moral residue – psychological, normative, usually both – remains.

Importantly, excuses make a difference to an action's psychological and normative residue without making a difference to its moral status – they do not negate its status as a wrong. This sets excuses apart from justifications, which bear on an action's moral residue by showing that it was not a wrong all-things-considered.

This invites a first objection. Why accept that moral remainders indicate moral responsibility? You might point to Williams' lorry driver to suggest that they may simply indicate bad moral luck. The driver ought to feel sorrow for his role in the accident that killed a child. But since the accident came about through no fault of his, he is not morally responsible for the death of the child. His connection to the death is merely causal.

But consider Williams' own discussion of the case. He says:

We feel sorry for the driver, but that sentiment co-exists with, indeed presupposes, that there is something special about his relation to this happening, something which cannot merely be eliminated by the consideration that it was not his fault. It may be still more so in cases where agency is fuller than in such an accident, though still involuntary through ignorance.<sup>49</sup>

The objection only succeeds if causal responsibility is enough to ground this "special relation". But I don't think it is. Many individuals are causally responsible for the death of the child: the lorry driver's employer who sent him on this particular route, the lorry manufacturer who produced the lorry. Yet while all of those various parties may regret the death of the child, it is only apt for the driver to feel sorrow and only he is an apt target for our pity. Why? Because while plenty of parties may have causally contributed to the death, it is only the lorry driver who *killed* the child. The significance of this is not causal but moral: it makes him, and not the others, morally responsible.

How could he be morally responsible, you might object, if he has not done anything wrong? Here we have to be careful. The lorry driver was neither reckless nor negligent: his brakes were checked, his speed below the speed limit, he was alert and watchful. Still, it is false that he has done nothing wrong: he *killed a child*. The killing was neither intentional,

---

<sup>49</sup> Williams [1981]. p. 28.

reckless, nor negligent. But our moral obligation is not to kill, not merely to refrain from killing intentionally, recklessly, or negligently.

Of course there's a difference between the lorry driver and a reckless road rowdy who kills a child in a deadly collision. The lorry driver has an excuse: the killing was unintentional (and neither negligent nor reckless). The reckless driver ought to feel remorse (indeed, be consumed by it), we are entitled to condemn and resent him. The unlucky driver, in contrast, ought to feel sorrow and grief and deserves our sympathy and pity. But it is only because he is morally responsible for killing the child that there is moral residue to contend with. The excuse does not negate the moral residue, it merely changes its nature. Thus, it would be wholly appropriate for the lorry driver to offer the bereaved family not only his condolences but also his deepest apology and to ask for their forgiveness. And as Matthew Kramer argues:

An apology accepts responsibility for a certain state of affairs or chain of events, and admits that that state of affairs or chain of events is wrong. It admits...that one's treatment of somebody else has fallen short of minimally acceptable standards. Though the falling short may be heavily extenuated by one's having adopted all reasonable precautions or by one's having had to avert an even worse outcome, it remains a falling short. As such, it calls for a moral remedy or a set of moral remedies. An apology concedes as much, without denying the possible existence and weightiness of some mitigating considerations.<sup>50</sup>

Williams' lorry driver is fictional but the analysis here chimes with the first-hand experience of those who have been involved in such fatal accidents. It is striking that these individuals self-attribute moral responsibility even when they are clear-eyed about the tragic circumstances and bad luck involved. A woman who, momentarily blinded by the evening sun, hit and killed a motorcyclist on her commute sums up:

"Yes, it was an accident, and in a certain sense we were both to blame, but, at the end of the day, I *hit* him, I took *his* life," she said. "No matter how much you want to dismiss it as an accident, I still feel responsible for it, and I am."<sup>51</sup>

I suggest that we should take such testimony at face value. In doing so, we heed Williams' advice who cautions us that...

...it would be a kind of insanity never to experience sentiments of this kind towards anyone [namely, the sorrow of the accidental killer and our pity for him], and it would be an insane concept of rationality which insisted that a rational person never would.

---

<sup>50</sup> Kramer [2005], p. 330-31. Kramer defends the view that moral liability is "strict" and that having taken impeccable precautions mitigate one's remedial obligations, rather than absolving one from them altogether.

<sup>51</sup> Gregory [2017], online at <https://www.newyorker.com/magazine/2017/09/18/the-sorrow-and-the-shame-of-the-accidental-killer>

To take the testimony at face value is to accept that the woman feels morally responsible because she really is morally responsible, rather than because she is conceptually confused or irrational. This is not to deny the moral luck involved. On the contrary, it allows us to diagnose its nature: the bad luck consists in her having become morally responsible for killing another human, through no fault of her own. As Williams observes about Oedipus:

“The terrible thing that happened to him, through no fault of his own, was that he did those things.”

One way of putting the claim that all excuses leave moral residue is to say that all excuses are *partial*. This may give you pause: surely, we don't treat all instances of excused wrongdoing alike. The agent who assaults another party after repeated provocation in a bar elicits a very different set of responses from the unlucky lorry driver.

But to say that all excuses are partial is not to say that all instances of excused wrongdoing are *on a par*. It is just to say that all excuses leave some moral residue. The nature and extent of this moral residue can vary widely, depending on the excused wrongdoing in question, its consequences (both moral and material), the circumstances in which it was committed, the relationship in whose context it took place, and the nature of the excuse itself: the wrongdoer may owe guilt, remorse, shame, sorrow, or regret, the wronged party may be entitled to feelings of anger, resentment, betrayal, sadness, or disappointment. Her reparative duties may range from the duty to make a quick apology to the duty to publicly admit guilt and pay compensation. I doubt there is a simple formula to determine moral residue – it will have to be determined case by case, by substantive normative inquiry.

This sheds light on the intuitive distinction between the lorry driver and the provoked bar assailant. In the case of the lorry driver, the circumstances of the accident negate our entitlement to resentment and outrage and his obligation to feel remorse. In the case of the bar assailant, the fact that she was provoked only mitigates these entitlements and obligations. Philosophical discussion of blameworthiness often focusses on resentment and outrage: to be blameworthy is to be a legitimate target for those reactive attitudes.<sup>52</sup> This can make it natural to say that the the lorry driver is not blameworthy, while the bar assailant remains somewhat blameworthy (although less than she would have been without the provocation). But we must not conflate blameworthiness, so understood, with moral responsibility. The lorry driver is not *less* responsible than the bar assailant; he is merely morally responsible in a different way. Both agents' excuses are partial; it's only with respect to resentment that the bar assailant's is more partial than the other.<sup>53</sup>

You may resist the thought that *all* excuses are partial. What about

---

<sup>52</sup> See e.g. Wallace [1994], p. 76

<sup>53</sup> I'm thankful to an anonymous reviewer for pressing me on this point.

psychosis, sleep walking, hypnosis, or an epileptic seizure? Surely those are full excuses: they leave no moral residue and fully negate moral responsibility.

I agree that hypnosis, psychosis, or being asleep absolve the agent from moral responsibility. But it is a mistake to classify these considerations as *excuses*. Rather they are *exemptions*.

There are principled reasons for distinguishing exemptions from excuses.<sup>54</sup> Excuses, as we saw, acknowledge that you have committed a wrong but mitigate your blame. But the fact that you have spilled the orange juice as a result of an epileptic seizure, hurt someone in the grip of a psychotic delusion, or stepped on someone's foot while sleep walking do not show that you have acted wrongly but are not responsible for it. Rather, they undercut the charge that you have *acted wrongly* in the first place.

They do so because to act wrongly you must violate a moral obligation that applies to you. But epilepsy, sleep walking, or a psychotic episode show that moral obligations simply do not apply to the agent in question at the relevant time. We do not expect a psychotic agent to keep a secret, nor do we expect a sleep-walker to watch where they are stepping. For this reason, a psychotic agent may harm you in giving away your secret but this harm does not constitute a wrong. Similarly, an agent undergoing an epileptic seizure may harm you by crashing into you but she does not wrong you in doing so. Exemptions tell us that the agent permanently or temporarily lacked the basic faculties that constitute moral agency – a grip on reality, knowledge of right and wrong, the capacity for self-control.

The question of excuses arises only after we have already established that the agent has acted wrongly: that is, that she is not subject to an exemption. Moral obligations do bind the fatigued, the stressed, and the distracted. To say that their circumstances excuse is not to cease holding the agent up to the relevant norms and expectations; it's just to recognize that her failure to comply with them is compatible with her having had a morally adequate present-directed intention.<sup>55</sup>

---

<sup>54</sup> The distinction between exemption and excuse is drawn by Strawson [1962]: “[exemptions] invite us to view the agent himself in a different light from the light in which we should normally view one who has acted as he has acted.” While I agree that exemptions show the agent in a different light, they also show the *action* itself in a different light: they show it not to be an instance of *wrongdoing*. For more discussion on the distinction between excuses and exemption, see e.g. Gardner [2007], p. 132. Wallace [1994], chapter 6 also distinguishes between excuses and exemptions along the lines discussed here.

<sup>55</sup> While excuses and exemptions are distinct normative phenomena, some considerations can, depending on context, function as either, for example “being just a child”. This can indicate that a certain moral obligation just does not apply – we do not expect a three year old to keep a secret. But it can also function as an excuse: “she was acting on a good intention but didn't foresee the consequences of her action – she is just a child.”

## 7. Conclusion

What are excuses? What makes a given consideration an excuse? I have argued that excuses are considerations that acknowledge that the agent acted wrongly but deny that the wrongdoing resulted from a lack of morally adequate intention. This account of excuses, the Good Intention Account, captures the wide range of excuses that we recognize in our moral practice in a unified framework and distinguishes excuses from the related normative phenomena of justifications and exemptions. What do excuses do? I have argued that excuses are responsibility-modifiers. They do not negate moral responsibility but modify the way in which someone is morally responsible for what they have done.

## 8. Work Cited

Allais, Lucy (2008). Wiping the slate clean: the heart of forgiveness. *Philosophy & Public Affairs*, 36(1), 33-68.

Arpaly, Nomy (2003). *Unprincipled virtue: An inquiry into moral agency*. Oxford University Press.

Arpaly, Nomy & Schroeder, Timothy (2014). *In praise of desire*. Oxford University Press.

Austin, John L. (1957) "A plea for excuses." *Proceedings of the Aristotelian Society*, New Series, Vol. 57 (1956 - 1957), pp. 1-30

Baron, Marcia (2007). "Excuses, excuses." *Criminal Law and Philosophy* 1.1: 21-39.

Baron, Marcia (2014). II—Culpability, Excuse, and the 'Ill Will' Condition. *Aristotelian Society Supplementary Volume* 88 (1):91-109.

Becker, Lawrence C., & Becker, Charlotte B. (2011). *Encyclopedia of ethics*. 2nd edition, Routledge.

Brandt, Robert. B. (1969). A utilitarian theory of excuses. *The Philosophical Review*, 337-361.

Bratman, Michael (1987). *Intention, Plans, and Practical Reason*. Center for the Study of Language and Information.

Calhoun, Cheshire (1989). "Responsibility and reproach." *Ethics* 99, no. 2: 389-406.

Duncan, John, Emslie, Hazel, Williams, Phyllis, Johnson, Roger, & Freer, Charles. (1996). Intelligence and the frontal lobe: The organization of goal-directed behavior. *Cognitive psychology*, 30(3), 257-303.

Friese, Malte, Loschelder, David, Gieseler Karolin, Frankenbach Julius, and Inzlicht, Michael (2018). "Is ego depletion real? An analysis of arguments." *Personality and Social Psychology Review*.

Gardner, John (2007). *Offences and Defences: Selected Essays in the Philosophy of Criminal Law*. Oxford University Press.

Gregory, Alice. (2017). The Sorrow and the Shame of the Accidental Killer. *The New Yorker*, September 18, 2017 Issue (<https://www.newyorker.com/magazine/2017/09/18/the-sorrow-and-the-shame-of-the-accidental-killer>)

Harman, Gilbert. (1986). *Change in View* (Cambridge: M.I.T. Press).

Harman, Elizabeth (2011). "Does moral ignorance exculpate?." *Ratio* 24.4: 443-468.

Harman, Elizabeth (2015). The Irrelevance of Moral Uncertainty. *Oxford Studies in Metaethics*.

Herman, Barbara (1993). *The Practice of Moral Judgment*. (Cambridge, MA: Harvard University Press).

Hieronymi, Pamela (2001). Articulating an uncompromising forgiveness. *Philosophy and Phenomenological Research* 62 (3):529-555.

Holton, Richard. (2009). *Willing, wanting, waiting*. Oxford University Press.

Holton, Richard & Shute, Stephen (2007). Self-control in the modern provocation defence. *Oxford Journal of Legal Studies*, 49-73.

Hume, David (1738/1960). *A Treatise of Human Nature*, ed. L. A. Selby-Bigge, Oxford University Press.

Hume, David (1748/1955). *An Inquiry Concerning Human Understanding*, ed. Charles W. Hendel, Bobbs-Merrill.

Inzlicht, Michael, & Milyavskaya, Marina (2017). Attentional and motivational mechanisms of self-control. In *Routledge International Handbook of Self-Control in Health and Well-Being* (pp. 11-23). Routledge.

Kelly, Erin (2013). "What is an Excuse?" In D. Justin Coates & Neal A. Tognazzini (eds.), *Blame: Its Nature and Norms*. Oxford University Press.

Kramer, Matthew H. (2005). Moral Rights and the Limits of the Ought–Implies–Can Principle: Why Impeccable Precautions are No Excuse. *Inquiry*, 48:4, 307–355

Markovits, Julia (2010). Acting for the Right Reason. *Philosophical Review*, 119(2), 201–242.

Moody-Adams, Michele M. (1994). Culture, responsibility, and affected ignorance. *Ethics* 104 (2):291–309.

Muraven, Mark, & Baumeister, Roy F. (2000). Self-regulation and depletion of limited resources: Does self-control resemble a muscle?. *Psychological bulletin*, 126(2), 247.

Murphy, J. G. (1982). Forgiveness and resentment. *Midwest Studies in Philosophy*, 7(1), 503–516.

Owens, David (2012). *Shaping the normative landscape*. Oxford University Press.

Rosen, Gideon (2008). "Kleinbart the oblivious and other tales of ignorance and responsibility." *The Journal of Philosophy* 105.10: 591–610.

Rosen, Gideon (2004). Skepticism about moral responsibility. *Philosophical Perspectives* 18 (1):295–313.

Rosen, Gideon. (2014). I—Culpability and Duress: A Case Study. In: *Aristotelian Society Supplementary Volume* (Vol. 88, No. 1, pp. 69–90).

Smith, Angela (2015). Responsibility as Answerability. *Inquiry* 58 (2): 99–126.

Smith, Michael (1994). *The Moral Problem*. Blackwells.

Strawson, Peter F. (1962). Freedom and resentment. *Proceedings of the British Academy* 48:1–25.

Tadros, Victor (2005). *Criminal Responsibility*. Oxford University Press.

Wallace, Jay R. (1994). *Responsibility and the moral sentiments*. Harvard University Press.

Williams, Bernard (1981). *Moral luck: philosophical papers 1973–1980*. Cambridge University Press.

Williams, Bernard (2008). *Shame and necessity* (Vol. 57). University of California Press.

Williamson, Timothy (forthcoming). Justifications, Excuses, and Skeptical Scenarios. In Dutant, Julien and Dorsch, Fabian (eds.). *The New Evil Demon*. (Oxford: Oxford University Press)

Zimmerman, Michael J. (1997). Moral responsibility and ignorance. *Ethics* 107 (3):410-426.

Zimmerman, Michael J. (1988). *An Essay on Moral Responsibility*. Rowman & Littlefield.