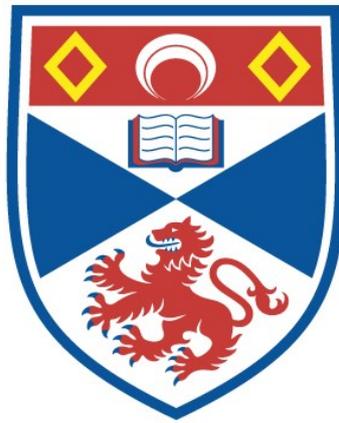


"YOU WANT ME TO DO WHAT?!" : A REASONABLE  
RESPONSE TO OVERLY DEMANDING MORAL  
THEORIES

Joe Slater

A Thesis Submitted for the Degree of PhD  
at the  
University of St Andrews



2018

Full metadata for this item is available in  
St Andrews Research Repository  
at:

<http://research-repository.st-andrews.ac.uk/>

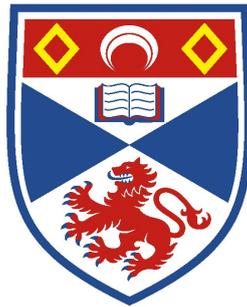
Please use this identifier to cite or link to this item:

<http://hdl.handle.net/10023/15546>

This item is protected by original copyright

‘YOU WANT ME TO DO WHAT?!’:  
A REASONABLE RESPONSE TO OVERLY  
DEMANDING MORAL THEORIES

Joe Slater



University of  
St Andrews

This thesis is submitted in partial fulfilment for the degree of PhD  
at the  
University of St Andrews

Date of Submission  
26th January 2018

# Declaration

## 1. Candidate's declarations:

I, Joe Slater, hereby certify that this thesis, which is approximately 79,000 words in length, has been written by me, and that it is the record of work carried out by me, or principally by myself in collaboration with others as acknowledged, and that it has not been submitted in any previous application for a higher degree.

I was admitted as a research student in June 2012 and as a candidate for the degree of Doctor of Philosophy in January 2014; the higher study for which this is a record was carried out in the University of St Andrews between 2014 and 2018.

Date                      signature of candidate:

## 2. Supervisor's declaration:

I hereby certify that the candidate has fulfilled the conditions of the Resolution and Regulations appropriate for the degree of Doctor of Philosophy in the University of St Andrews and that the candidate is qualified to submit this thesis in application for that degree.

Date                      signature of supervisor:

## 3. Permission for publication:

In submitting this thesis to the University of St Andrews I understand that I am giving permission for it to be made available for use in accordance with the regulations of the University Library for the time being in force, subject to any copyright vested in the work not being affected thereby. I also understand that the title and the abstract will be published, and that a copy of the work may be made and supplied to any bona fide library or

research worker, that my thesis will be electronically accessible for personal or research use unless exempt by award of an embargo as requested below, and that the library has the right to migrate my thesis into new electronic forms as required to ensure continued access to the thesis. I have obtained any third-party copyright permissions that may be required in order to allow such access and migration, or have requested the appropriate embargo below.

The following is an agreed request by candidate and supervisor regarding the publication of this thesis: No embargo on print copy. No embargo on electronic copy.

Date                      signature of candidate:

signature of supervisor:

# Abstract

This thesis is about demandingness objections. It is claimed that various moral theories ask too much of moral agents, and for that reason should be rejected or modified accordingly. In the first chapter, I consider what this objection entails, particularly distinguishing it from Bernard Williams's integrity objection. The second chapter investigates several attempts to undermine the objection. I contend that their arguments for a more burdensome conception of morality fail, and that accepting their 'extreme' view would leave us unable to explain much of our moral phenomenology. In the third chapter, I analyse what features of a moral theory make it susceptible to demandingness objections. Through this discussion I highlight social factors (the conduct and expectations of one's community) and psychological factors as potential candidates for generating the problem. Making use of these potential diagnoses, in chapter four, I examine (but ultimately reject) the responses to demandingness objections by Richard Miller and Liam Murphy, which can provide verdicts sensitive to these features.

In the fifth chapter, I examine the concept of blame and its relationship to moral wrongness. Noting this relationship and how an action's difficulty can affect whether we deem conduct blameworthy, I consider a recent proposal by Brian McElwee, that the difficulty of certain actions explains why they are too demanding. I reject this proposal, instead regarding difficulty as providing excuse conditions. However, through the discussion I draw attention to the fact that sub-optimal behaviour often does not need an excuse, suggesting that there is no 'default' obligation to do the best. In the final chapter, I offer a way to consider how obligations are generated, utilising the concept of reasonableness. By incorporating this concept, and giving it a relativistic analysis, I suggest a theory can avoid demandingness objections.

# Acknowledgements

I am extremely grateful to many people who have helped make this thesis possible. My supervisors, Justin Snedegar and Tim Mulgan, have proved invaluable. Without their assistance, which has been supererogatory, what follows would be greatly impoverished or, more likely, unfinished. This project first came to my attention through conversations and presentations by Brian McElwee in St Andrews in 2011. Since then, Brian has become both my friend and a huge influence on my philosophical development.

I am also thankful to Sarah Broadie, John Haldane and Ben Sachs for constructive feedback whilst reviewing my progress. I also benefited considerably from conversations with Theron Pummer, particularly about chapters 2 and 5. Some of the discussion in the final chapter draws on discussions with Marcia Baron, who has been kind enough to send me and allow me to use unpublished materials. The St Andrews graduate seminar series has also been very helpful in providing an engaging audience and insightful questions throughout the last few years.

The fine establishment of Aikman's, the St Andrews Philosophy Society and the University's Shinty teams must also be credited for offering me regular and much needed escape from the perils of Thesis-Land. I'd also like to express my gratitude to the many people who write funny things on Twitter for entertaining me during the work-day.

The title of this thesis was the source of a great deal of controversy in the final few weeks. The final title arrived at was settled on thanks to suggestions by Josh Habgood-Coote and Farbod Akhlaghi-Ghaffarokh. Though I didn't opt for their suggestions, thanks for alternatives provided by Lea Wisken, Lisa Bastian, Jonathan Ichikawa, Ste Broadrick, Janis Schaab and Joshua Seigal. My thesis defence itself was a surprisingly pleasant experience. For that I must thank my examiners, Brad Hooker and Theron Pummer, for the positive and enthusiastic manner they conducted the viva.

My sanity throughout the process of writing this owes a lot to the music of Regina Spektor. I definitely couldn't have pulled off many of the seemingly endless all-nighters in the office without music playing constantly (and vast quantities of tea), and Regina Spektor's albums were often just what I needed.

Finally, I would like to thank the St Andrews Philosophy Department

in general. Aside from the fact that they were the source of my funding for this PhD, the encouragement (and toleration) they have offered me over the years has been amazing. I have lived in the town for over eleven years now. Had the department not been so welcoming, accessible and friendly over this time, I would not still be here, nor be the person I am today.

# Contents

<b>1</b>	<b>What are Demandingness Objections?</b>	<b>4</b>
1.1	Different Objections . . . . .	7
1.1.1	“Standard” Demandingness Objections . . . . .	13
1.1.2	Williams’s ‘Integrity’ Objection . . . . .	19
1.1.3	Epistemic Demandingness . . . . .	26
1.2	Different Objections and Different Types of Responses . . . . .	29
1.2.1	Agent-centred Prerogatives . . . . .	30
1.3	Overview . . . . .	33
<b>2</b>	<b>Is the Objection Successful? Responding to the Extremists</b>	<b>34</b>
2.1	Sobel and the ‘Impotence’ of Demandingness Objections . . . . .	35
2.2	Singer . . . . .	40
2.2.1	Singer’s Arguments for Strong Demandingness . . . . .	40
2.2.2	Rejecting Singer’s Argument . . . . .	44
2.3	Unger . . . . .	52
2.3.1	Peter Unger’s Argument for Demandingness . . . . .	52
2.3.2	Rejecting Unger’s Argument . . . . .	59
2.4	Dealing with Extremists . . . . .	65
2.4.1	Sidgwick . . . . .	66
2.4.2	Denying Overridingness . . . . .	68
2.4.3	Evaluation . . . . .	70
<b>3</b>	<b>What Makes a Theory Overly Demanding?</b>	<b>71</b>
3.0.1	Aside: Types of Cases . . . . .	71
3.1	Obligation-Generation . . . . .	73
3.2	Impartiality . . . . .	80
3.3	Insensitivity to Social Factors . . . . .	86
3.4	Insensitivity to Individual Psychology . . . . .	92
3.5	The Diagnosis . . . . .	93
3.6	Conclusions . . . . .	94

<b>4</b>	<b>Potential Solutions and Failings</b>	<b>96</b>
4.1	Miller's Underlying Dispositions . . . . .	97
4.1.1	Problems of Miller's Account . . . . .	100
4.1.2	Rejecting Miller's view . . . . .	107
4.2	Murphy's Collective Beneficence . . . . .	108
4.2.1	Problems with Murphy's View . . . . .	114
4.3	Explaining the failures . . . . .	120
<b>5</b>	<b>Blameworthiness, Difficulty and Obligation</b>	<b>121</b>
5.1	Blameworthiness and Moral Wrongness . . . . .	123
5.1.1	Some Thoughts About Blame . . . . .	123
5.1.2	The Relationship Between Blameworthiness and Wrong- ness . . . . .	126
5.1.3	Blame and Demandingness . . . . .	132
5.2	Difficulty . . . . .	134
5.2.1	An Argument from Blame . . . . .	139
5.2.2	Difficulty as Merely an Excuse Condition . . . . .	143
5.2.3	Resistance to Understanding Difficulty as Merely an Excuse Condition . . . . .	148
5.2.4	Rejecting Difficulty and Reconceptualising the Problem	149
5.3	Foundations of Moral Obligation: Success Criteria for a Moral Theory . . . . .	154
<b>6</b>	<b>Positive Proposal: Reasonableness in Obligation</b>	<b>157</b>
6.1	Views of Reasonableness . . . . .	162
6.1.1	Sibley and Rescher . . . . .	162
6.1.2	Rawls . . . . .	163
6.1.3	Scanlon . . . . .	165
6.1.4	Marcia Baron on Reasonableness . . . . .	167
6.1.5	General features of reasonableness . . . . .	169
6.2	Reasonableness in Theory: A Very Brief Sketch . . . . .	171
6.3	Reasonableness in Theory: Scanlon's Contractualism . . . . .	174
6.3.1	How Demanding is Scanlon's Contractualism? . . . . .	174
6.3.2	Problems in Scanlon's Contractualism . . . . .	179
6.3.3	Why this matters for Demandingness . . . . .	188
6.4	Relative Reasonableness . . . . .	193
6.4.1	Conclusions so far . . . . .	197
6.4.2	Other Issues of Demandingness . . . . .	198
6.4.3	Relativistic Problems . . . . .	203
6.4.4	Success Evaluation . . . . .	207
6.4.5	A Concession to the Extremists . . . . .	211

# Introduction

In 1972, Peter Singer's "Famine, Affluence and Morality" issued a challenge to both common-sense moral thought and the way philosophers regard obligations to the global poor. Following on from a very simple principle, that "If it is in our power to prevent something very bad from happening, without thereby sacrificing anything of comparable moral importance, we ought, morally, to do it" (Singer, 1972, p.231), Singer argues that our obligations are extensive. This principle has the virtue of sounding, *prima facie*, extremely plausible. It becomes clear, however, after some consideration that the principle inevitably (in any world with such rife amounts of easily preventable suffering, like ours) entails that we all ought, morally, to devote enormous levels of our time/money/energies to preventing suffering.

The vast amount of obligations generated is easily seen when considering typical everyday choices. When deciding what to do for dinner, perhaps I think it would be nice to try the new restaurant, but then I realise that I could buy some pasta and cook it myself for much less money, and donate the rest to a charity, which could use the difference to prevent a child going blind. It seems undeniable that a child becoming blind is "something very bad". It also seems obvious that my sacrificing trying the new restaurant in favour of a cheap meal at home is not comparably morally significant. This procedure can then be applied to any luxury purchase one might make.

We are, if we are to accept Singer's conclusion, committed to an austere lifestyle, devoid of luxuries except in situations where indulging in them will not forestall our ability to relieve some of the widespread suffering. Many find the conclusion that morality requires this much of us to be absurd; this would just be too demanding. Responses to arguments and wider theories like this, which take issue with the claim that morality is *that* demanding, have come to be known as demandingness objections.

While current discussions of demandingness often begin with Singer, this is not a new problem. Sidgwick discussed the accusation that utilitarianism set up “too high a standard of unselfishness” and that it made “exaggerated demands on human nature” (1907, p.87). Mill, also defending utilitarianism, noted the complaint that its standards are “too high for humanity” (1861, II.19).<sup>1</sup> In previous centuries, most people did not have the disposable income, the requisite knowledge or the ability to assist those suffering the most. However, technological developments have increased our ability to help the distant poor, and the world’s wealthy (even those who are modestly well-off in developed nations) have amassed what would previously have been regarded as tremendous wealth.

Though historically consequentialist moral theories have been the primary target for these types of objections,<sup>2</sup> there have been several discussions recently of how demandingness could be problematic for other moral theories. Demandingness concerns are now raised in relation to Kantian ethics,<sup>3</sup> Scanlon’s contractualism,<sup>4</sup> “common-sense morality”<sup>5</sup> and virtue ethics.<sup>6</sup> The question of how demanding morality is therefore provides a huge problem for moral philosophers.

Objections of this sort are also not merely problems for academic philosophers. Considerations of overdemandingness permeate many discussions relating to many aspects of human life. Fiona Woollard, for example, notes that demands typically expressed in the media of new or expecting mothers are excessively demanding, and identifies this as problematic (2016b). There are also worries that structural issues in society impose excessive demands on care workers.<sup>7</sup> Overdemandingness is also noted, though not in the same vocabulary, in popular media with regards to a variety of issues. There are concerns regarding the demands of being an ethical consumer – trying to purchase products that do not unnecessarily harm animals or impact the

---

<sup>1</sup>References to *Utilitarianism* will be given in the form “A:n” where ‘A’ is the chapter and ‘n’ is the paragraph.

<sup>2</sup>Exactly why this is the case is a difficult and interesting question. It is my contention that the relative simplicity of these theories as well as the clarity of their verdicts are likely explanations, though I do not address these hypotheses in this thesis.

<sup>3</sup>e.g. van Ackeren and Sticker (2015), Pinheiro Walla (2015), Timmermann (forthcoming).

<sup>4</sup>e.g. Ashford (2003), Hills (2010), Ashford and Mulgan (2012, §8).

<sup>5</sup>e.g. Fishkin (1982).

<sup>6</sup>e.g. Swanton (2009).

<sup>7</sup>e.g. Tessman (2015).

environment – which is often taken to require too much of people.<sup>8</sup> The ideals of beauty which face women also present another structurally similar objection, with the expectations seeming to ask too much.<sup>9</sup>

In this thesis I aim to explore what demandingness objections actually involve, what makes a theory vulnerable to this type of objection, and how we should respond to overly demanding theories. In the opening chapter, I consider several different objections that risk being conflated and all relate somehow to demandingness. In the second chapter, I examine a variety of arguments which attempt to blunt the force of demandingness objections, ultimately rejecting these efforts. In chapter 3, I investigate several candidates for features a moral theory could have which might result in overdemandingness. This proves to be difficult because none of the obvious candidates seem necessary for making demandingness objections. In the 4th chapter, I look at some specific attempts to respond to the objection in ways that looked promising considering the findings of chapter 3, but eventually reject these attempts. In chapter 5, I take seriously recent attempts to explain overdemandingness in terms of difficulty, and though I reject these attempts, suggest that exposing why they fail may elucidate some of our problems. Finally, I offer a tentative schema of a moral theory based on reasonableness, which I suggest can avoid accusations of overdemandingness.

---

<sup>8</sup>e.g. Rohwetter (2017)

<sup>9</sup>e.g. Hinsliff (2017).

# 1

## What are Demandingness Objections?

“That’s too much, man.”

– Sabrina, *Horsing Around*<sup>1</sup>

Demandingness objections are typically forceful when a moral theory, principle or practice requires a great deal from individual moral agents. The obvious candidate for a demanding moral theory, and a regular target of the objection, is act utilitarianism. It is well-noted in the literature that this was the original target of the objection.<sup>2</sup> Act utilitarianism holds that the best action one can take in any given situation is the one that results in the greatest utility, or most happiness. When combined with the requirement to maximise, as is often how act utilitarianism is regarded,<sup>3</sup> this becomes an extremely demanding theory.

We can see how demanding maximising act utilitarianism is by considering the obligations one would ordinarily be subject to. Kagan depicts this issue succinctly in *The Limits of Morality*:

“When I go to the movies I may spend a few dollars and enjoy myself for an hour or two. The pleasure I get is genuine, and it

---

<sup>1</sup>Fictional show featured in *Bojack Horseman*

<sup>2</sup>For examples, see Cullity (2009, p.8), Hooker (2009a, p.151),

<sup>3</sup>Though this is often treated as the standard (as can be seen in many instances where the maximisation requirement goes unspecified, as in Murphy (1993) and Hooker (1991)) and is actually endorsed by several utilitarians (e.g. Sidgwick (1907), Singer and de Lazari-Radek (e.g. 2016)), this is by no means necessary to act utilitarianism or consequentialist positions in general, as noted by Norcross (1997, p.2) and McElwee (2007).

seems absurd to say that I have done anything wrong. Yet this is exactly what the claim [maximising act utilitarianism] entails, for both my time and my money could be better spent: the pleasure one could bring in an evening visiting the elderly or the sick quite outweighs the mild entertainment I find in the movies; and the money itself would have done much more good were it sent to famine relief – for even a few dollars is sufficient to enable another human being to survive a temporary food shortage brought on by drought. If the claim is right, then in going to the movies I do what is morally *forbidden*.”

(Kagan, 1989, p.1)

For any state of the world that resembles the way the world is now – when there are people starving and dying of easily preventable diseases – maximising act utilitarianism commits agents to extreme demands. It seems unlikely that one could be justified in purchasing any luxury goods,<sup>4</sup> <sup>5</sup> so every agent is likely to be morally compelled to have a very austere life. Demandingness objections hold that any moral theory which requires this of moral agents in situations like our own must be mistaken, that “there are limits to the demands of morality’s requirement that we promote the good” (Murphy, 1993, p.274).

It should be noted that demandingness objections do not simply reject a theory because it may lead to extremely high demands (though this is compatible with everything said up to this point). We could call such an objection the *Primitive Demandingness Objection* (PDO). To clarify, I envisage this sort of objection as proceeding as follows:

P1: No moral theory should require more than X of an agent in any situation.

---

<sup>4</sup>An exception to this is when the purchase of luxury goods can be expected to increase one’s ability to help the world’s poor, by giving one more influence over wealthy persons or enhancing one’s career prospects, which could be instrumental in helping the world’s poor. Singer makes this point in *Practical Ethics* (1993, p.223,n.1).

<sup>5</sup>Throughout this thesis I assume that the affluent westerners are not culpable for the extreme suffering in the developing world. As Judith Lichtenberg points out, our complicity in harmful global systems makes this dubious (2010), but as taking this kind of consideration into account would ‘muddy the ethical waters’, I discuss these matters as if this was not an issue.

P2: Moral theory A demands more than X in some situations.

C: Moral theory A must be rejected.

\*Where X denotes some degree of sacrifice.

With the tacit premise that moral theories that demand more than they should must be rejected, this argument is clearly valid, but it still seems both problematic and unlikely to properly encapsulate the difficulty identified by people like Murphy and Kagan.

Presumably we accept that there are situations wherein extreme demands are legitimately made of agents. In some extreme situations we would expect agents to be faced with extreme demands, and it might even be a fault of a theory if it did not deem that in those cases agents are subject to very challenging moral obligations.

To illustrate this, Ashford notes that a “moral theory must hold that there are some situations in which agents face extreme moral demands – for example, a situation in which the only way of stopping billions of people suffering an agonizing death was by hacking off your left leg with a fairly blunt machete” (Ashford, 2003, p.274). We clearly see this as an extremely demanding requirement, but given the situation, we are still likely to deem this morally required. If the legitimacy of this demand is questioned, we may look for far more ordinary cases where morality makes extreme demands of us. If one becomes a parent – and has not relinquished care to some other responsible party – it is expected that if the situation makes it necessary, the parent will take extremely demanding steps to ensure the wellbeing of the child. This is likely to amount to a considerable amount of time and effort being expended over the course of many years, a significant financial cost and, in many cases, the sacrifice of several personal projects one may identify with.

Similarly, any theory which holds that the act of legitimate promising confers moral obligation upon the promiser must allow that such promises could be very demanding. An agent can promise to go to extreme lengths, and if we think the promise is legitimate, it would ordinarily be expected that the agent would be under a moral obligation to fulfil this commitment.

Something of the form PDO could still be held and allow that these demands are legitimate, but only if X is held to be at an extremely high level, so that none of these obligations would be above that level. Perhaps,

requiring a person to end her life would be a point where we might suppose the demand cannot be legitimate.<sup>6</sup> We might take issue with this in two ways.

- (i) We might suggest, and quite plausibly so, that there are situations which will morally require one to sacrifice one's life.
- (ii) We might suggest that theories prescribing demands which fall far short of this may be objected to because of their demandingness.

The second of these seems the more pressing (as the first could be responded to by simply adjusting the X-value even higher). It seems that demandingness objections, properly understood, can be levelled against theories that require significantly less in some cases than sacrificing one's life. To illustrate this, let us consider a revised version Ashford's machete case. Instead of cutting off one's leg with a machete to save billions from an agonizing death, doing so would simply cause a brief but significant pleasure to a group of a thousand.<sup>7</sup> If this is claimed to be a moral requirement under a given moral theory, a demandingness objection may be applicable. The lesson to be learned here is that the amount that can be legitimately demanded of an agent is sensitive to the situation they are in.

Accepting that demandingness objections are not simply concerned with how much is demanded of an agent by a moral theory, we might wonder exactly what the objection is. Many commenters in the demandingness literature seem to be making use of different ideas of what the objection is, which may be quite problematic. In the remainder of this chapter I examine several types of objection that may be conflated.

## 1.1 Different types of Objection

Before briefly examining several objections from the literature, it is worth introducing some terminology which is useful in describing some features

---

<sup>6</sup>Hobbes made a claim of this sort; that while moral obligations come from the state, the sovereign cannot legitimately demand that a citizen forfeits her life (e.g. *Leviathan*, 1995, XIV).

<sup>7</sup>Under a simplistic consequentialism, we can manipulate the numbers enough to make the consequences better if the agent does perform the act. There are other objections that might be made to this sort of case, such as aggregation related issues (e.g. (Scanlon, 1998, p.220-240)), but a demandingness objection is also available.

of demanding moral theories. In *Human Morality*, Scheffler discusses four types of responses to moral theories which make “unusually heavy demands on individual moral agents” (Scheffler, 1992, p.17). First, it might be denied that morality is as forceful as the theory suggests; that the demands are too burdensome and that a less taxing theory should be adopted. This response denies that morality is as stringent – as able to make demands of such a magnitude – as the theory suggests. This type of response, it might be said, denies *stringency*.

Second, we might claim that the scope of morality is narrower than the theory claims. There might be cases where it seems inappropriate for morality to provide a verdict for an action. It would, for example, sound slightly strange to claim that it is morally permissible to brush one’s teeth (Scheffler, 1992, p.20). Another type of example where it might be claimed that moral justifications or explanations are inappropriate occur when an agent is faced with a choice where one option is so costly it seems unthinkable. Scheffler gives an example of a man choosing between saving the life of his wife or a stranger.<sup>8</sup> The decision here might seem to lie “beyond justifications”.<sup>9</sup> This type of case might be more readily associated with demandingness objections, because of the heavy demands that could arise from treating one’s wife’s life equally to the stranger’s here. According to some theories (e.g. Taurek, 1977), a coin should be flipped (or some other fair decision procedure taken) to make the decision, which does seem to ask too much of an agent, in a sense. This type of response, which claims a moral theory is overly moralistic, denies *pervasiveness*.

Thirdly, one might accept that morality does make the purported burdensome demands, but deny that an agent is compelled, all-things-considered, to acquiesce to these demands. For example, we might believe that what one ought to do morally is distinct from one ought to do rationally. As this sort of response rejects the notion that what one morally ought to do always trumps any other considerations,<sup>10</sup> this denies *overridingness*.

The final response Scheffler sees available to a heavily demanding moral

---

<sup>8</sup>Scheffler borrows this example from Bernard Williams’s *Moral Luck* (1981, p.18).

<sup>9</sup>This phrase is used by both Scheffler (1992, p.19) and Williams (1981, p.18).

<sup>10</sup>Dale Dorsey endorses a view of this sort, (e.g. “The Supererogatory and How to Accommodate it”, 2013), which I will consider in 2.4.2. For now, I simply want to claim that though demandingness objections can be made whatever position one takes on overridingness, if one believes morality is rationally overriding, these objections may hold additional force.

theory is to simply accept the burdensome demands. Several authors have suggested that because of the state of the world, morality should be demanding, and that we are required to act in accordance with it. This response accepts stringency, pervasiveness and overridingness, but at the cost of violating our common intuitions about the demands of morality.<sup>11</sup>

It seems clear that though denying the pervasiveness or the overridingness of morality are reactions to the demands levelled by a moral theory, these are not objections to the *demandingness*. If one objects to the pervasiveness of a moral theory, the objection is not because of how demanding a purported obligation is, but that there should not be any obligation at all. The complaint concerns the subject matter of the demand, not its demandingness. An objection about the pervasiveness suggests that a category error has been committed, as morality has no business making verdicts in whatever domain is in question.

Denials of morality's overridingness are also not concerns of demandingness of a moral theory, though these concerns are more likely to have a bearing, albeit indirectly, upon how demanding we think a moral theory should be. If we deny that morality is overriding we may say that morally, an agent should  $\phi$ , while accepting that all-things-considered, it is permissible not to  $\phi$ . According to this view, an agent may be morally blameworthy for not  $\phi$ -ing, but not rationally criticisable. Those who hold such a view may not find demandingness objections as convincing, because the failing to act in accordance with moral obligation will *only* be a failing in one particular domain of human activity, whereas for others (like Kant, for example), it would be a failure of practical rationality. I will leave this aside however, for the time being, as it is certainly possible to have demandingness objections whatever one's position on overridingness.

Even if we restrict our discussions of demandingness to problems concerning stringency, there are still (or so I shall claim) several ways the objection can be interpreted.

In recent years demandingness objections have received a reasonable amount of attention. Since Peter Singer's "Famine, Affluence and Morality" made a widely-appreciated case for affluent people being morally required to do much more for the global poor than they currently are, many have

---

<sup>11</sup>Responses of this general type are offered by Singer and Unger, and are discussed in 2.2.1 and 2.3.1 respectively.

discussed whether and how we should amend our moral theories in response.

Despite a great deal of attention to demandingness objections in recent years, I suggest that it is not obvious what the purported problem actually is. Often a short example is given to demonstrate what the problem is supposed to be, or instances of the problem. To illustrate this, I will give some examples from prominent figures who have discussed the demandingness objection.<sup>12</sup>

When some authors describe the problem, they take pains to emphasise the importance of personal projects. Tim Mulgan, who has written extensively on consequentialism, and authored one book entitled *The Demands of Consequentialism*, writes: "...the common objection that Consequentialism is unreasonably demanding, as it leaves the agent too little room (time, resources, energy) for her own projects or interests. . . I shall call this the Demandingness Objection" (2001, p.4). Here we might note that Mulgan locates the problem in the restriction of one's personal projects (or interests).<sup>13</sup> Similarly, David Sobel suggests that the problem the objection purports to identify is that morality "properly understood, should not take over our lives, at least in circumstances such as we face these days, but should be compatible with a range of attractive and self-directed lives, including lives that involve a serious commitment to family, friends, or non-moral projects" (2007, p.3). Again we may notice an emphasis on personal projects. I discuss why it might be problematic to regard the objection in this fashion in section 1.1.1.

Garrett Cullity suggests that there are actually three problems of demandingness. He is a pluralist about morality, and he sees moral demands deriving from (at least) three sources: a morality of concern, a morality of respect and a morality of cooperation. Cullity suggests that each of these has a corresponding demandingness problem (2009, p.12). The familiar problem of demandingness pertains to concern and is usually introduced in terms of

---

<sup>12</sup>For present purposes I will simply assume that those talking about "the demandingness objection", "the problem of moral demandingness" or "the over-demandingness objection" are all talking of the same issue. It might be that certain authors have picked their terminology specifically to differentiate themselves from others in the literature, but as they make reference to each other, without (as far as I have seen) specifying that they have a different notion in mind, and thus might be talking past each other, I assume this is not the case.

<sup>13</sup>As is mentioned later (in 1.1.1) Mulgan does give a more detailed analysis (2001, p.25). That analysis does not mention personal projects, but merely 'unreasonable demands'.

welfarist theories: “If what morality requires of us is the impartial promotion of welfare... it also seems to imply that almost all forms of pursuit of my own interests are wrong” (2009, p.13). This seems to clash with the standard thought that any “morality we can have reason to accept must be one that allows us to lead full and fulfilling lives – lives that are open to and directed towards the goods that make for a flourishing human existence” (2009, p.13). This conflict brings out the problem in terms of the morality of concern.

With regards to the “morality of respect”, Cullity suggests a parallel problem arises from respecting the autonomy of others. If an agent was obligated to act in a way that did not curtail any of the positive or negative freedoms of others to any degree, Cullity suggests we may suffer extreme demands. From respecting the positive right of others to pursue a career, for example, we might be forced never to compete with another for a job (2004, p.16) as this may result in a restriction of their autonomous plans. Similarly, Cullity suggests that in the morality of cooperation, if an agent took part in all the cooperative projects one has good reason to, this would also “leave no time for ourselves” (2004, p.19).<sup>14</sup>

Liam Murphy, who has also written extensively on demandingness, claims “the ‘over-demandingness objection’, as I will call it, asserts that there is a limit to how great a sacrifice morality, or at least a principle of beneficence, can legitimately demand of agents” (1993, p.268). This suggests that there are certain sacrifices that cannot be obligatory; anything which is more demanding than the “limit” Murphy mentions.

Brad Hooker, in his chapter in Chappell’s collection of essays concerning demandingness, describes the objection fairly pithily, stating that “the objection levelled at act-utilitarianism [the demandingness objection] is that act-utilitarianism demands more sacrifice than it is reasonable to demand” (2009a, p.151). This characterisation seems unobjectionable, but perhaps not particularly illuminating. Later, he goes on to claim “the demandingness objection to act-consequentialism is that it intuitively seems that morality

---

<sup>14</sup>I will not properly discuss the idea that there are different objections relating to a pluralist conception of the source of moral demands. One reason for this is that Cullity’s pluralistic analysis of the roots of demandingness seems to be a minority position. Additionally, all the problems Cullity identifies will be accounted for in the general form of the “standard demandingness objection”, which is given in the following section and will go on to be my primary focus.

doesn't demand making contributions constantly for the sake of aggregate welfare impartially assessed, or for the sake of aggregate welfare impartially assessed plus equality, or for the sake of aggregate welfare assessed in a prioritarian way" (2009a, p.153). This is certainly more specific, but it is an analysis specifically catered to act-consequentialism. Hooker discusses demandingness with regards to act-consequentialism and some variations upon it – and his formulation of rule consequentialism – in several places<sup>15</sup> but does not offer a more general analysis.

Brian McElwee, in a categorisation of different types of objections that may be conflated with demandingness objections, defines demandingness objections as the claim that an . . .

“ . . . agent is morally obliged to do A should be rejected just because doing A is too costly to the agent, given the moral considerations at stake, to count as morally obligatory.”

(McElwee, 2017, p.93)<sup>16</sup>

This exposition holds that the type of objection we are looking for applies when a moral theory identifies some action as obligatory, but the action requires too much (time/energy/sacrifice) of the agent. McElwee notes that while any theory for which this description applies will be too demanding, sometimes a more fundamental objection will apply. He considers a theory which requires agents “to spend Saturday afternoon assaulting as many innocent strangers as [they] can” (2017, p.87). While a verdict like this will often be too demanding (I tend to think I am required to do very little on Saturday afternoons!) it seems like there is a more obvious problem with this theory; going round hitting people isn't even *good*, or permissible, let alone obligatory. To distinguish cases like this from instances where the demandingness of the theory is the appropriate complaint, McElwee describes “pure demandingness” objections as obtaining when a theory states that “doing A is what is morally best, and what is overall best, but is not morally obligatory because the cost to the agent is too great” (2017, p.90). This *pure*

---

<sup>15</sup>e.g. “Rule-Consequentialism” (1990), “Rule-Consequentialism and Demandingness: A Reply to Carson” (1991), *Ideal Code, Real World* (2000).

<sup>16</sup>This is the analysis I most favour among the literature, though this should come with the caveat that the ‘A’ in question may be some disjunctive list of actions. It seems clear that a theory may be overly demanding while permitting choice between several options, all of which are too demanding.

objection targets situations wherein a theory correctly identifies the best action to perform, but incorrectly (or so the objection holds) requires this optimal behaviour. For instance, we could accept that act utilitarianism provides the correct evaluation of which act is morally best, but dispute that the act is thus morally obligatory, and do so purely because the cost to the moral agent is so high.

Sophie-Grace Chappell importantly notes that “the problem of demandingness is not the problem that morality, or a moral theory can make extreme demands.”<sup>17</sup> It is important to be clear that the objection at hand is not of this sort because there are situations, like if “a million lives can only be saved by my self-sacrifice” (2007, p.255) when such demands can legitimately be made.

Though I think Chappell is correct in this negative claim, her analysis seems a little strange in other ways. She claims that “the problem of demandingness is not about a theory’s making extreme demands sometimes. The problem is that some moral theories seem to make extreme moral demands on us, not just sometimes, but always” (2007, p.255). Chappell suggests that the problem lies in the constant intrusion into our daily lives imposed by a moral theory, which “threaten important parts of human life” (2007, p.256). In the following section, I will suggest that this is problematic, and provide my own analysis of the problem of moral demandingness.

From the variety of comments above we can note that some authors have highlighted the role of personal projects in demandingness objections. In the next section, I offer my own analysis of the objection, showing, among other things, why I think an emphasis on personal projects might be mistaken.

### 1.1.1 “Standard” Demandingness Objections

Despite all that I have said so far regarding what a demandingness objection is (or isn’t), I do not regard it as a complicated type of objection. It requires that a moral theory issues demands upon moral agents. It also needs these demands to *ask too much* – to require inappropriately extensive sacrifice from the moral agents. For an obligation to be inappropriately demanding (in the way that may generate a demandingness objection), there must be some amount (of time/energy/sacrifice) that is appropriate to demand, and

---

<sup>17</sup>This particular point is in line with the rejection of PDO as a way of understanding the demandingness objection (detailed in the opening section of the chapter).

the obligation must demand more than this.

I would suggest that in most instances the *appropriate degree* is given to us by pre-theoretic intuition. We just *have a feeling* that doing a certain amount is morally required – and that more is not. That the appropriate amount required is given to us in this way is supported by the way the objection is often presented. As Murphy puts it, the stringency of the demands “strike just about everyone as absurd”, and prompts us to think that these verdicts “just couldn’t be right” (Murphy, 2000, p.6). This does not rely upon any particular moral epistemology, only that we come to these beliefs about what is (and is not) required of us independently of theoretic concerns. This pre-theoretic intuition is usually likely to be coarse-grained, so not always able to provide verdicts of specific actions. We would not expect, for example, these intuitions to give exact quantities one should donate to charity every month, but might expect judgments of non-obligation for particularly high proposed figures, or that more is required than any proposed low figures. Though I think this is the way the purported appropriate degree of demandingness usually features, I would not want to rule out any other method of determining how this might be calculated.

What is important is that the appropriate degree is sensitive to different situations. We might imagine utopian societies in the future where minimal or no moral demands are appropriate. In that sort of situation, a modest requirement (e.g. for a moderately well-off person to relinquish 5% of their income or a few hours a week) might seem extraordinarily demanding. Alternatively, we can consider situations of an opposite sort, rife with extreme suffering, where it is appropriate for extreme demands to be made of an agent.

One further concern is that we presumably want to be able to apply a demandingness objection to theories even if the overly demanding situations do not actually (or currently) obtain. If a moral theory is only moderately (and acceptably) demanding now, but is just as demanding in a utopian society, an agent might still rightly object to the moral theory, simply because it would make inappropriately demanding obligation verdicts in a nicer world. We typically accept that moral theories should not just provide us with verdicts in the current state of affairs, but also make correct predictions about hypothetical situations.

For this reason, I would reject Chappell’s claim that “the problem is

that some moral theories seem to make extreme moral demands on us, not just sometimes, but always” (2007, p.255).<sup>18</sup> Even if the world were one day to approach a utopian ideal, wherein a maximising utilitarianism would be minimally demanding we would still think that the theory was subject to demandingness objections, simply because the theory *would* be so demanding in less favourable circumstances.

With these considerations in mind, I propose the following general formulation of the demandingness objection (TDO):

D1 We have some reason<sup>19</sup> to believe, in a specific situation (or set of situations),<sup>20</sup> that a moral agent is subject to moral obligations of a certain degree of demandingness.

D2 In these situations (or sets of situations), a moral theory/principle gives a verdict which demands significantly more than this.

C We have reason to believe that this moral theory should be rejected.<sup>21</sup>

22

---

<sup>18</sup>We might accept this claim if it is restricted to the world today. For most relatively affluent people in the western world, the extreme demands we face (according to maximising consequentialism, at least) are ubiquitous. Chappell may have intended this qualification, but it is not made explicit.

<sup>19</sup>This might seem like a particularly weak claim, as we have *some reason* to believe all sorts of things. The strength of a specific objection will depend, however, upon how good the reason is. A weak reason to believe obligations are of some weight will make a weak objection.

<sup>20</sup>Demandingness objections often operate from multiple situations because the extreme demands are created by iterating many seemingly small demands. These are paradigmatic instances where the objection takes hold (e.g. Kagan (1989), Hooker (1999), Mulgan (2001), Bourbaki (2001), Cullity (2004)).

<sup>21</sup>The strength of the conclusion will depend on how strong our reasons are in the first premise.

<sup>22</sup>This evaluation is similar to the form provided by Mulgan when considering the objection specifically towards consequentialism (2001, p.25). Mulgan analyses it as:

- “1. Consequentialism makes demand D;
2. D is an unreasonable demand for a moral theory to make; therefore,
3. Consequentialism makes unreasonable demands”

This would be one instance of the analysis I have provided, where the moral theory in question is consequentialism and the reason to believe the degree of demandingness one is subject to (or not subject to) is its unreasonableness. My formulation is adapted slightly to allow for mistaken demandingness objections (where, though it might appear to the contrary, morality *really is* that demanding) and to allow for a variety of reasons for thinking that a theory is excessively demanding.

While TDO is *available* in any cases that meet this criteria, it will not always be *appropriate*. If the theory in question requires actions that are morally irrelevant or even impermissible (like McElwee’s Saturday afternoon example), other types of objections will be more suitable. A demandingness objection is particularly *appropriate* when the verdict provided in D2 correctly identifies an action which is morally better than the alternatives (or a set of actions morally better than the alternatives).

### **An Aside on Projects**

It might be noted that my analysis, provided above, makes no mention of personal projects. Some might find this surprising, considering how personal projects featured prominently in the descriptions of the objection by Mulgan, Sobel and (as will be discussed later) Miller.<sup>23</sup> I hope to demonstrate, however, that the sacrifice of personal projects is neither necessary nor sufficient for a demandingness objection to be available.

Consider an objection that is based upon the sacrifice of personal projects, which I will term Project Demandingness (PD):

- P1: Morality should allow a person to pursue personal projects to a significant degree.
- P2: Moral theory X does not allow the pursuit of personal projects to this degree.
- C: Moral theory X should be rejected.

What exactly “personal projects”<sup>24</sup> is taken to mean might be somewhat difficult to define, but for the time being, we can simply accept that these are non-moral activities<sup>25</sup> which enrich the lives of individuals.

If the demandingness objection is taken to be of this structure then it must be the case: i) that every instance of a demandingness objection

---

<sup>23</sup>Miller devises his own principle – the principle of *Sympathy* – in response to Singer, for the express purpose of allowing agents to partake in more personal projects (2004, p.359). This is discussed in section 4.1

<sup>24</sup>As will be noted later, Bernard Williams talks extensively about personal projects, as well as *commitments*, which he views as a special type of project.

<sup>25</sup>By this I mean activities undertaken where moral concerns are not a constitutive part of the agent’s intention — though they are likely to cause the agent happiness, which will make them instrumentally good under many moral theories.

involves the restriction of one's personal projects, and ii) that every moral theory that does not allow pursuit of personal projects to a sufficient degree is vulnerable to this objection. I dispute both of these claims, arguing rather that a restriction of one's personal projects is neither necessary nor sufficient for a demandingness objection to be made against a moral theory.

Let us first consider when a demandingness objection to a moral theory should exist. A theory must issue some demands. It isn't the demands themselves – that is, the content of the demands – that will be problematic however, but *how demanding* they are. The same requirement might well differ in how demanding it is for different moral agents to satisfy. It may be demanding in terms of how difficult it is to fulfil, or how much an agent may need to sacrifice (time/energy/options/proportion of resources). To constitute an objection, we must also stipulate that the amount demanded is inappropriate, otherwise, the demandingness of the requirements would be no problem for the theory.

With these considerations in mind, we can now consider two ways PD may fail to describe demandingness objections. On one hand we may examine whether there are demandingness objections which may be made in cases requiring no sacrifice of an agent's personal projects. On the other, we might see whether we may find cases where a moral theory may demand a significant sacrifice of one's personal projects where a demandingness objection is not available. Of the first type, let us consider **Noble Susan**.

**Noble Susan:** Susan is a remarkable person. She is extremely selfless and all her personal projects are extremely noble, to the extent that the best acts she may perform in any situation involve pursuit of these projects. She cares a great deal about eradicating poverty, for example. A maximising form of consequentialism is likely to deem her morally obliged to spend inordinate amounts of time, and almost all her resources, on relieving the suffering of the world's poor. Though she may be truly committed to this cause, it would seem mistaken to say she was morally obliged to do this much. Here then, we appear to have an agent never having to sacrifice any of her personal projects, but demanded by the moral account to do something extraordinary (what we would most likely view as supererogatory).<sup>26</sup>

---

<sup>26</sup>A case of this general sort is described briefly by Scheffler, who describes someone whose project “simply *was* to bring about the best state of affairs” (1982, p.22).

This type of case does seem to lend itself towards a demandingness objection. The requirements purportedly placed on Susan seem inappropriate even if it is the case that she wants to do these things. Of the second type of case we may consider **Zombie Apocalypse**:

**Zombie Apocalypse:** Fred is a gardening enthusiast. Growing different plants, reading about gardening and looking at or smelling plants are his favourite things to do. He spends over 80% of his waking life doing gardening-related activities, and when he sleeps, he usually dreams of gardening. It is safe to say that gardening is Fred's top priority personal project. Unfortunately, one day there is a zombie apocalypse! Zombies are everywhere, and Fred knows the only way his two children will be able to survive is to hide in his basement for a considerable amount of time (enough to let all his favourite plants die!). He will have to sacrifice all his gardening-related fun for the foreseeable future, but we probably think this is morally required of him.

Applying PD to this case, it seems that Fred should be morally permitted to engage in his gardening, and thus that any theory prohibiting that should be rejected. However, it seems obvious that giving up his gardening endeavours is required in this situation. It is clear that most moral theories require Fred to sacrifice his projects here, and that this requirement is not overly demanding. Thus, there are cases wherein a theory may require that an agent relinquish one's projects and such a theory need not be overly demanding. We do not have to seek examples which are quite so extravagant for cases of this sort. If we think one can acquire legitimate moral obligations by promising, and these obligations may prohibit pursuit of personal projects, this would be sufficient.

If arguments of these types are accepted as demonstrating either a demandingness objection or a Project Demandingness (but not both), we must accept that the sacrifice of personal projects is neither necessary nor sufficient for a demandingness objection. It is obviously the case that in many situations, according to moral theories which seem inappropriately demanding, an agent will also have to sacrifice personal projects. However, as I hope to have demonstrated in this section, the connection is not essential to demandingness objections.

I propose that one explanation for the role of personal projects in the

discussion of demandingness objections is the influence of Bernard Williams, which I discuss in the following section.

### 1.1.2 Williams's 'Integrity' Objection

Bernard Williams is regularly noted as having brought attention to a difficulty relating to demandingness for utilitarianism. In *Utilitarianism: for and against* he makes complaints about utilitarianism on two major grounds; its difficulties in catering adequately for justice, and the notion of integrity (1973, p.82). It is with regards to the latter that Williams's objections concern (or are purported to concern) demandingness.

Leading up to his discussion of integrity, Williams provides two cases he believes utilitarianism gives strange verdicts for. In the case of George (1973, p.97), a chemist with few job opportunities is presented with the possibility of working for a laboratory conducting research into chemical and biological warfare. Though George is reluctant, he is informed that if he does not take the job, someone who is extremely enthusiastic about developing weapons and lacking in scruples will get it. Thus, it seems George could prevent more powerful weapons being created and further deaths resulting by taking the job.

In the second case, a botanist named Jim has accidentally ended up in a strange South American village during an expedition (1973, p.98). Twenty Indians are tied to a wall. Jim is informed that they had protested against the government and were to be killed to dissuade further protesters. Because of Jim's arrival as an honoured guest to their land, however, he is offered the chance to kill just one of the Indians himself. If he does so, he is assured that the other nineteen will be released to mark the occasion. If he does not, the killing of all twenty will go ahead as planned.

Because utilitarianism, as presented by Williams, ignores any special importance given to one's particular role in causing a state of affairs, the verdicts utilitarians arrive at in both of these cases are supposed to be extremely obvious. George should take the job, and Jim should kill the Indian. Because the states of affairs brought about by these acts are much better than the alternatives, utilitarianism suggests that the correct choice is so clear that no serious deliberation need be undertaken. The fact that it would be Jim *doing* the killing does not enter as a consideration, and neither does George's *doing* the development of the biological weapons. Williams

holds this as an important reason why utilitarianism is wrong; these factors *should* come into consideration. Williams concedes that in the case of Jim and the Indians, the verdict is “probably right” (1973, p.117). The issue is not with the ultimate verdict, but with the deliberation process and the agent’s relation to it.

Utilitarianism, as broadly construed by Williams, leaves many open questions, which an advocate of the theory may develop in various ways. The commitments he does ascribe to it are that it is a form of “consequentialism” and “eudaimonism” and that it is “direct” (1973, p.79). Eudaimonism is the view that what is desirable (and all that is desirable) in itself is happiness. Williams leaves open exactly how this should be thought of; some may suggest pleasure, others satisfaction, or some other particular notion. He uses the “economists’ phrase” utility to pick out that property being valued by the theory. A theory’s being ‘direct’ entails that the concern of morality is individual actions, rather than rules or practices. (It might be noted that this does mean that rule utilitarianism, by Williams’s account, is not classified as a form of utilitarianism.)

The notion of consequentialism, it turns out, is much more complicated than might at first be supposed. As Williams is eager to point out, “no one can hold that everything, of whatever category, that has value, has it in virtue of its consequences” (1973, p.82). This, of course, would lead to a vicious regress. We must then accept that something has value in virtue of something other than its consequences. An obvious option here is to consider states of affairs as bearers of value/desirability, the degree to which they are valued/desirable depending on the sort of consequentialism we are considering (so, for utilitarianism, happiness).

Williams introduces the notion of the *right* action, which he sees as important to consequentialism.<sup>27</sup> He holds this notion to involve maximising for consequentialists; the right action will be that which results in the best consequences (again, according to the criteria of the particular theory). This will commit a utilitarian to the following claims:

---

<sup>27</sup>This is a claim I would dispute. For versions of consequentialism that allow many options (ones which do not require maximising), though there may be a *best action*, this is not the *right action*. We may consider scalar approaches, where there are no right or wrong actions (e.g. Norcross, 1997), or satisficing alternatives where above some standard there will be many permissible actions (e.g. Slote 1984). The notion of *the* right action implies that all other possible actions are wrong (so presumably forbidden).

- I. In a given situation  $S$ , the moral agent did the right thing by doing  $A$ .  
*entails*
- II. The state of affairs  $P$  is better than any other state of affairs accessible to her.

Where  $P$  is a state of affairs caused by her doing  $A$ .

This entailment distinguishes consequentialism from alternative moral theories. Kantian moral theory, for example, would deny this entailment. The categorical imperative would forbid breaking promises, for example, even if doing so would lead to a state of affairs wherein fewer people broke promises.

What an agent has an obligation to do thus depends on what state of affairs is brought about by her actions. The notion of the state of affairs brought about by the action here is one that includes the action, and what utility that involves, including the utility gained by the agent (though this only counts as much as the utility gained by any other agent would). We can note now how *negative responsibility* plays such a big role. It does not matter whether a state of affairs consists in what an agent does, is produced by what the agent does, is allowed by the agent, encouraged by the agent, etc. (1973, p.93). Whether the bringing about of a state of affairs requires that an agent acts or omits to act does not matter.

Whatever the criterion the form of consequentialism judges as valuable, like promotion of happiness for utilitarians, it must value some projects, other than the general project of bringing about maximally desirable states of affairs. Among these might be first order projects: “desires for things for oneself, one’s family, one’s friends, including basic necessities of life. . . , objects of taste” as well as “pursuits and interests of an intellectual, cultural or creative character” (1973, p.110). Without valuing some first-order projects, at least indirectly, a form of consequentialism “would have nothing to work on and would be vacuous” (1973, p.110).

In our everyday lives, we also value second-order projects (those which involve other projects). Williams suggests that in order to be at all palatable, any form of consequentialism must also value such projects (and not simply the general utilitarian second-order project of maximising first-order projects). Among such projects would be those that “flow from some more general disposition towards human conduct and character, such as hatred of

injustice, or of cruelty or killing” (1973, p.111). Projects of this category are described as *commitments*, and we may be involved in many, like a “cause, an institution, a career, one’s own genius, or the pursuit of danger” (1973, p.112).

Even if a utilitarian does account for commitments (by rightfully recognising that these are sources of happiness), they still fail, according to Williams to do so appropriately. As their commitments can only affect the evaluation of their actions as one unprivileged set of commitments among many, the agent may be alienated from their own commitments. *Their* commitments play no special role in moral deliberation. This no longer treats the agent as a *person*, but rather “just a representative of the satisfaction system who happens to be near certain causal levers at a certain time” (1981, p.4). It is this lack of a special role for individual moral agents’ projects and commitments that Williams sees as problematic for utilitarianism. He sees this as an affront to individual integrity.

### Exposing the Differences

It might already appear clear, given what has been said in the previous subsection, that the objection to utilitarianism presented by Williams given above, is distinct from TDO.<sup>28</sup> In case that is not the case, here I shall attempt to delineate the objections, as well as give some explanation as to why the objections may be conflated. This is important because (as is discussed in section 1.2) these objections may warrant different responses, and some authors do regard Williams’s objection as a demandingness objection. Chappell, for example, describes Williams’s integrity objection as “his version of the demandingness objection” (2007, p.255).<sup>29</sup>

---

<sup>28</sup>That Williams’s objection is different to the demandingness objection has been noted before. Conly distinguishes between two ways the demands of a moral system can be “dangerous”; they can demand too much sacrifice (loosely, the demandingness objection) or they may *alienate* a moral agent (1985, pp.78-281). Railton also discusses a similar objection to Williams under the name of alienation (1984). Mulgan also notes that these objections are logically distinct (2001, p.16).

<sup>29</sup>Braddock, similarly sees Williams’s objection as a demandingness objection (2013, p.169). Ashford also seems to regard the integrity objection as related to the demandingness objection, but seemingly *because* loss of integrity is such a significant demand (2000). Liam Murphy regards Williams’s objection as a *combination* of both the demandingness objection and the alienation objection: He claims there is “first, the claim that these theories are too demanding: it is too frequently morally impermissible to pursue projects and relationships. And then there is the second claim that to the extent that we are permitted to pursue our projects and relationships, impartial morality alienates us from

As what has been discussed so far – with particular emphasis on the decision procedure and how it fails to take the individual and their projects into account – seems quite different to what is usually referred to by demandingness literature, that this objection is seen as related to demandingness might seem peculiar. One relation may be found in Williams’s discussion of integrity. He claims that, when thinking of someone like Jim in the case of the Indians, it...

“... is absurd to demand of such a man, when the sums come in from the utility network which the projects of others have in part determined, that he should just step aside from his own project and decision and acknowledge the decision which utilitarian calculation requires.”

(1973, p.116)

So here we have a demand made on him by utilitarianism, which is unreasonable – *absurd* even. Even here though, we must note that it is not the verdict (or how demanding it is) which is absurd. As Williams concedes that Jim’s killing one Indian is (probably) the right decision, it can’t be the verdict – that Jim should kill the Indian – that is wrong here. It is the *alienation* of an agent from their own actions and projects, which utilitarianism (in Williams’s account) requires that he cites as problematic.

One reason we may take pains to distinguish demandingness objections from alienation objections (which I do concede provide a difficult challenge) is that we seem to be able to pull these objections apart.

### **Demandingness without alienation:**

We might think that the case of Noble Susan (from section 1.1.1) is one where a demandingness objection is available, but which does not feature alienation from one’s commitments.<sup>30</sup> However, it must be remembered that the alienation objection focusses on the lack of consideration for an agent’s own projects qua *their* projects. Even here the agent’s projects are offered

---

them” (1993). I take Williams’s claim to be more of the second sort, and disentangling that is the focus of this section.

<sup>30</sup>The goal of alleviation of suffering — featured in the Nobel Susan case — would certainly count as the type of personal project which is also a commitment under Williams’s terminology.

no special place in her thoughts. It may be considered fortunate for her that the best action available to her happens to align with her commitments, but she is still hostage to utilitarian calculation. Susan could still complain that her projects being hers have no relevance to the calculation.

In order to provide an example of a demandingness objection where there is no alienation objection, something more extravagant might be needed. We might consider an agent without any commitments. It must be remembered that *commitments* are those projects “with which one is more deeply and extensively involved and identified” (1973, p.116), so this would not be the same as an agent being without preferences or tastes. Still, such an agent would probably be a bit strange; one accumulates commitments as a natural part of life. To illustrate something of this sort, take Soap Opera Sally:

**Soap Opera Sally:** Sally suffers from a soap-opera style amnesia (where she does not remember who she is or anything about herself, but has a normal grasp of language and everyday activities). Sally has no impetus to discover what she was like before. After a chance encounter with Peter Singer in the recovery room, and proceeding through her utilitarian calculation, she discovers that she is required to commit vast energies to relieve the world’s poor.

This still seems like too much is being demanded, and that a demandingness objection can be made. Clearly she is not being alienated from any of her commitments, however, as she has none.

#### **Alienation without demandingness:**

It also seems that we can create cases wherein an agent is alienated from her commitments, but a demandingness objection is not available. One might suggest that the cases of Jim and George could be such cases, but as it might be thought by some that the obligations on these are too demanding, I will try to provide a clearer case.

**Lucky (?) Lucy:** Imagine Lucy, who is in the (admittedly rare) situation wherein doing exactly what she would have wanted to do is the optimal action. Perhaps she is living in some utopian society which makes no or minimal demands, or has found herself in a situation where, for

the time being, she cannot affect the lives of anyone else in any significant way. She performs her calculations, and results after consulting the utility network inform her that she should do precisely what she wanted to do prior to her calculations.

This seems like a good situation for Lucy, but if the decision is given unto the agent via the various complicated sums from the “utility network”, not because of any special relationship she has to the type of action, alienation worries seem to persist. This may still, as Williams puts it, alienate her “in a real sense from [her] actions and the source of [her] action in [her] own convictions” (1973, p.116).

One might respond that the mere process of thinking about things impartially, not counting one’s own interests as special, is itself demanding. This does not seem convincing. Perhaps we may consider this situation, but wherein all the options available to one are fairly pleasant, and the calculation fairly simple. The calculation is not (by stipulation) epistemically demanding,<sup>31</sup> and thinking about the matter impartially hardly seems demanding. The process, however, may still divorce one from one’s own desires/projects/commitments, which is crux of the alienation objection.

As these examples have shown, it is difficult to pull these two objections apart because not considering one’s projects (or not considering them as having any special relationship to one’s actions) in decision procedures will strike one as demanding, and doing anything particularly demanding is likely to detract from one’s ability to partake in one’s personal commitments, and thus cause some alienation. This latter alienation however, need not arise from a lack of special consideration in one’s decision procedures, thus not be of the sort Williams is concerned with.

If what I have argued in the preceding sections is correct, it might be wondered why some writers on demandingness have understood Williams as making a demandingness objection. I suggest three answers to this question.

Firstly, Williams has clearly had an influence on many authors in the demandingness debate. Personal projects or commitments now feature prominently in most discussions of demandingness, and this may well originate in Williams’s comments on these matters.

---

<sup>31</sup>Epistemic demandingness and its role (or lack of) in relation to general moral demandingness will be discussed in section 1.1.3.

Secondly, several people have misunderstood the purpose of Williams's examples. Several people who discuss Jim and the Indians concern themselves with asking whether or not Jim *should* kill the Indians, and think the objection to utilitarianism lies in that it says he should.<sup>32</sup> If this was the type of objection Williams was making, then it could fit the form of a demandingness objection as described earlier. However, as noted in the previous section, that is not the case.

Finally, people may have been given the impression that Williams's objection is one of demandingness by a passage in *Utilitarianism: for and against* mentioned earlier. Williams writes, of moral agents like Jim and George, that it "is absurd to demand of such a man, when the sums come in from the utility network which the projects of others have in part determined, that he should just step aside from his own project and decision and acknowledge the decision which utilitarian calculation requires" (1973, p.116).

I suggest that this passage could be misleading as it explicitly mentions the demands faced by the agents, and also is phrased in a way that permits misunderstanding. Williams himself agrees that – at least in Jim's case – the agent should "step aside from his own project", but contends that he is not required to *just* do this, acknowledging and submitting to the utilitarian calculation, as this is to divorce him from his actions. This is clarified in the following sentence, where Williams asserts that to submit to utilitarian calculation as a disinterested party is "to alienate him in a real sense from his actions and the source of action in his own convictions" (1973, p.116).

### 1.1.3 Epistemic Demandingness

There are several issues which may be considered to pose problems for moral theories that fall under the general heading of 'epistemic demandingness'. These problems may obtain when a moral theory requires that agents go through a particular epistemic procedure or accept certain beliefs. I shall briefly discuss two of these, suggesting that they need not be considered at length here.

Firstly, we may consider a moral theory which is extremely cognitively

---

<sup>32</sup>In Chappell's analysis of Williams's integrity objection, she notes that R.M. Hare and Ted Honderich are among writers who have fallen victim to this particular mistake (Chappell, 2007, p.256, footnote 3).

demanding to actually live by.<sup>33</sup> A moral theory may be mentally strenuous to actually adopt, if for example it requires a large amount of calculation to discover what one ought to do (or may do, or may not do), taking so much energy that we would have to spend all our time calculating and perhaps still not be able to reach the appropriate verdicts. Mill mentions opponents of his who make this type of objection, who claim “there is not time, previous to action, for calculating and weighing the effects of any line of conduct on the general happiness” (1861, II,24). Mill does intend his theory to be practicable, so defends on those grounds, claiming that over the course of time we have had time to understand what acts promote happiness. Whether or not we agree is immaterial to the question at hand, as the discussion here is about the *correctness* of moral theories, not of how they may best be implemented. The issue Mill had in mind only acts as an objection to the claim that we can use the moral theory at hand in our day-to-day lives.

Scheffler also responds to this criticism, considering the accusation sometimes levelled against utilitarianism, that “utilitarian agents must always be monitoring their circumstances to make sure that no major opportunity” to promote the good slips by without their notice (1992, p.36). Utilitarians may respond to this in several ways, such as “appealing to the distinction between the utilitarian account of right action and the utilitarian account of optimal deliberation” (1992, p.37), which may differ. We could accept that a moral theory may be correct, but it would be better either to live as though it was not, or adopt a simplified decision procedure by which we may approximate it for practical purposes. Even act utilitarians often recommend thinking in terms of rules, because doing so will ensure more widespread understanding and/or result in better consequences (e.g. Singer and de Lazari-Radek 2010, p.55).

With this type of thought in mind, we can form two distinct objections:

1. That *there is no* decision procedure that could be appropriately action-guiding, i.e. that any calculation that could point towards right actions would require a greater more evidence and/or more computational capacity than we possess.
2. That there *is* some correct decision procedure the theory would en-

---

<sup>33</sup>Braddock notes that this is one interpretation of what could be meant by the claim that a theory “demands too much” (2013, p.170).

dorse, but that procedure is too demanding.

In the vein of the first type of manoeuvre, James Lenman considers the problem of foreseeing distant consequences. Not only, it seems, is the future we would bring about by specific actions *uncertain*, but we are actually *chueless* about the long-term effects. This raises the difficulty that a consequentialist about a theory of the good could not simply suggest that a different decision procedure would promote that good, as we are not in a position to even make reasonable predictions (2000, p.349). To illustrate this, Lenman considers the fictional example of Richard, who lived over two thousand years ago. Richard was a bandit, who raided towns and killed the inhabitants. In one such town, he left alive a woman called Angie. Angie, however, happened to be a distant ancestor of Hitler, so, it seems Richard's non-killing of Angie becomes a terrible wrong from a consequentialist position. Particularly with actions that determine whether someone lives or dies, our actions are obviously 'identity-affecting' – they change which persons live in the future. When these are considered, our actions clearly have massive causal ramifications.

If our actions have massive causal ramifications, which we are completely unable to predict, it is difficult to see how we would be able to properly evaluate the moral significance of any action (on a consequentialist picture). If that is the case, it is difficult for a consequentialist to then suggest any suitable decision procedure which is likely to result in the best (or even good) consequences. The difficulty for the consequentialist then becomes apparent, that in order to be in a good position to think their actions are good, they must perform what seems to be an impossible deliberative process.

This type of objection, from what Lenman calls the Epistemic Argument, highlights difficulties agents may have in arriving at an epistemic position that would allow them to promote the good. While there are questions that consequentialists must face here, it should be noted that these do not concern how demanding the theory is. Lenman's complaint is that consequentialism is incompatible with a decision procedure that can identify right actions (or even gesture towards better actions) rather than how demanding those actions might be.

For this reason, this is not really a problem of demandingness. It does not contend that the theory makes demands which are too hard, but rather

that we can not know what demands it would make.<sup>34</sup>

Now the second problem, that making decisions the way that a theory would recommend, even acknowledging that this might not require thinking in terms of the theory itself, might require more cerebral effort than we think is appropriate. If we consider this type of objection, however, it simply seems like a version of TDO given in section 1.1.1, but where what is demanded is going through a type of thought process.

The problems discussed in this subsection, thus pose either no real problem in terms of demandingness, or are easily subsumed under a more standard analysis.

## 1.2 Different Objections and Different Types of Responses

In the previous sections I have introduced several objections that could be made, all which relate to demandingness in some way. While it is a plausible reading of what some of the descriptions of overdemandingness, the ‘Primitive Demandingness Objection’ does not get off the ground, because it fails to appreciate that different situations may legitimately demand different amounts from an agent. ‘Project Demandingness’ similarly fails because it does not appreciate that some situations legitimately require one to abandon her projects.

Of objections that do provide problems for theories, I discussed TDO, Williams’s integrity objection and some concerns relating to epistemic requirements. As the epistemic problems seem either to not concern explicitly to a theory’s demandingness, or seem to be captured by TDO, these too can be disregarded for current purposes.

It is important to distinguish between the two remaining objections – TDO and Williams’s integrity objection – because the problems require different responses. To illustrate this we may consider the example of agent-centred prerogatives. I contend that utilising agent-centred prerogatives can provide a theory with the means to reply adequately to integrity objections, but not to demandingness objections.

---

<sup>34</sup>This is a serious worry for consequentialists, and one which a response is needed for, but this is not the question at hand.

### 1.2.1 Agent-centred Prerogatives

Samuel Scheffler made use of agent-centred prerogatives in order to respond to the view that there is discrepancy between the way consequentialism generates obligations – in a strictly impersonal way – and how “concerns and commitments are naturally generated and sustained from a person’s point of view quite independently of the weight of those concerns and commitments in an impersonal ranking of overall states of affairs” (1982, p.14). So agent-centred prerogatives would allow an agent to “assign a certain proportionately greater weight to his own interests than to the interests of other people” (1982, p.20).

As he later described it, agent-centred prerogatives mean that an agent is . . .

“... permitted to perform his preferred act (call it P), provided that there was no alternative A open to him, such that (1) A would produce a better overall outcome than P, as judged from an impersonal standpoint which gives equal weight to everyone’s interests, and (2) the total net loss to others of his doing P rather than A was more than M times as great as the net loss to him of doing A rather than P.”<sup>35</sup>

(1994, p.169)

This manoeuvre allows us to get the expected verdicts in several cases. For example, we might consider organ donation cases. If an agent is aware that someone else’s life could be saved (or significantly prolonged) by her donating a kidney or part of her liver with very low risk to herself, it looks like this would plausibly be the best thing for her to do.<sup>36</sup> However, most people do not see this as required. Agent-centred prerogatives can explain this by allowing the agent to assign a significantly greater weight to her own interest. Though donating the organ would produce the best outcome overall, the loss to the other person would be not greater than the loss to her

---

<sup>35</sup>Where ‘M’ is a multiplier indicating how much an agent is allowed to privilege one’s interests.

<sup>36</sup>Given that there are over 5,000 people on the waiting list for kidneys in the UK, and that 250 people on this list die every year (statistics from <http://www.giveakidney.org/why-we-need-more-donors/>, accessed January, 2018), most healthy people *are* in this situation.

(the risk of complications and the slightly lower quality of life which might be incurred) multiplied by the agent-centred weighting.

It should be noted that these prerogatives are *allowances*, rather than requirements. An agent is permitted to give her own interests or projects a greater weight when deliberating, but not required. They provide agents with options to act in a morally sub-optimal way. Agent-centred *restrictions*<sup>37</sup> on the other hand are restrictions which are “at least sometimes impermissible to violate in circumstances where a violation would prevent either more numerous violations, of no less weight from an impersonal point of view, of the very same restriction, or other events at least as objectionable, and would have no other morally relevant consequences” (1982, p.80). One example would be the Kantian restriction on lying,<sup>38</sup> as it is not deemed permissible to lie in that picture even if doing so would prevent more lies in the future. Scheffler rejects these restrictions, claiming they would be problematic, as they would commit an advocate to the position that “it is sometimes *impermissible* to produce the best state of affairs” (1982, p.22).

Scheffler reacts in this way because the conventional (maximising) consequentialism does not permit any actions which promote a non-optimal outcome, and supposes that this can be avoided if one’s own interests are given a greater weighting. Importantly, the prerogatives will allow agents to “permissibly pursue their own projects and commitments” but “not permit one to pursue one’s own projects at all costs” (1982, p.21).

Theories involving a use of agent-centred prerogatives purport to escape Williams’s integrity objection because by their very nature they allow for the special relationship between an agent and her projects. *Their* projects can (if the agent so chooses) count for more in their deliberations.

We might suspect that a similar tactic could work against TDO. If an agent is allowed to give their own interests special treatment, in a situation where they may have had to donate an extremely demanding amount, after a theory is amended with agent-centred prerogatives, this obligation would be reduced or eliminated. If we still thought the burden was too high, we could change the weighting until it seemed low enough.

However, Garrett Cullity points out that moves like this, which deny the

---

<sup>37</sup>Shelly Kagan refers to agent-centred prerogatives as “agent-centred options” and agent-centred restrictions as “agent-centred constraints” (e.g. (Kagan, 1989, ch.1)). I will retain Scheffler’s terminology.

<sup>38</sup>e.g. Kant’s “Of the Alleged Right to Lie from Philanthropy”.

impartiality of morality, might not be effective (2004, ch.6). Though they are successful in appropriately reducing some extremely demanding obligations (like the organ donation case discussed above), if we consider obligations that are very easy to discharge, it runs into a problem. Consider a situation wherein an agent can save a life (or perhaps many lives) at a negligible cost. Presumably we accept that the agent should act in this situation, or at the very least, that there are some situations like this where this is required. The explanation for this, according to the advocate of agent-centred prerogatives, must be that an agent is not permitted to assign their own interests a weight so much greater than the interests of others as to make it permissible not to act in such situations.

However, as we have seen, very burdensome demands often accrue by iterations. Donating the means to provide a malaria net for one child in the developing world may not be hugely demanding, but when there is another child, and another, these individually negligible obligations accumulate.

If an extreme demand is generated by an iteration of similarly small demands,<sup>39</sup> no plausible weighting to one's own preferences would prevent the high demandingness arising (2004, ch.6; 2009, p.14). Advocates of agent-centred prerogatives must take one of two options. Either they can place a weighting one assigns to one's own interests so high that the first instance of giving is not required or the iteration must continue until a situation obtains where the originally negligible amount could plausibly be enough, with the weighting, to excuse not donating. The first of these disjuncts — making the weighting extremely high so negligible amounts cannot be demanded — would rule out paradigmatic instances of moral obligation we would generally accept. Taking this route, assigning an agent's interests such an astronomically high weighting that no one else's interests could affect their obligations to others would commit one to denying any obligation in any of the seemingly obvious life-saving cases.

The second option, even with agent-centred prerogatives seems likely to lead to a Singer-esque conclusion, since the loss of the negligible amount (the cost of saving a life), in order to outweigh the saving of a life, must constitute a non-negligible cost to the agent. The problem is particularly

---

<sup>39</sup>Cullity accepts that one is obligated to contribute to group projects, and that these obligations persist even if others are not. As one's personal share in the group project (where the group is all affluent people) of saving *one life* may be extremely small (less than pennies), it is clear that this requirement would be negligible (2004, ch.5).

pertinent because it seems like this is the sort of situation which currently obtains in our everyday lives with regards to the world's poor.

Whereas it seems relatively easy to account for the integrity objection, by simply adding to a moral theory a variable to add additional consideration for an agent's commitments, no similar solution seems available for TDO.

### 1.3 Overview

In this chapter I have introduced the notion of demandingness objections and how they are thought to provide problems for various moral theories. I then argued that there were potential conflation about what was being discussed, and discussed different types of objection that might be considered. In section 1.2, I demonstrated that for the types of objection that might be thought to provide genuine difficulties for moral theories, different responses are applicable.

From this point forward, I focus on "Standard" objections.<sup>40</sup> One motivation for this is that this seems the one most deserving of the name – it is the one which specifically concerns demandingness and locates the problem there. When this objection is appropriate, it might be understood as a *pure* demandingness objection.

In the next chapter I consider whether the pure demandingness objection is successful – whether it successfully locates a problem with the theories it deems overly demanding.

---

<sup>40</sup>Henceforth, any discussion of "demandingness objections" refers to objections of this sort.

## 2

# Is the Objection Successful? Responding to the Extremists

*“The fault, dear Brutus, is not in our stars,  
But in ourselves”*

– William Shakespeare, *Julius Caesar*

As discussed earlier, there is a huge intuitive pull to reject a moral theory that seems too demanding. Consider a middle-income western person who donates (*only*) 30% of their income and (*only*) one night a week to volunteering. The suggestion that this person is morally failing strikes us as absurd.

A related argument for rejecting extremely demanding theories comes from the notion of supererogation, the idea that an agent can perform actions that are good but not required, or *beyond duty*. The usual examples of supererogation are instances of heroism, like the soldier who dives on a grenade to save others, or a doctor who risks their life working in a plague-stricken city.<sup>1</sup> According to theories like maximising act utilitarianism, this type of behaviour is required of anyone. Acting in that way is *merely* doing what one had to do. Viewing these types of actions that way seems to mischaracterise them considering how we feel about the agents performing

---

<sup>1</sup>These are examples given by Urmson in “Saints and Heroes” (1958, pp.202-203) but have become paradigmatic examples.

them. They do not seem to be doing ‘good enough’, but rather something incredible. Clearly, it seems, they did not *have to* do what they did. This line of thought seems like common sense, and if correct, seems to entail that there is something wrong with moral theories that claim people do *have to* act in that way.

An obvious response to demandingness objections is to deny that there is anything wrong about an extremely demanding moral theory. I have neglected this point up until now, but in this chapter I consider motivations for accepting that morality simply is extremely demanding (or at least it is when the world is like it is today). I consider the position of some who happily accept such a claim. Peter Singer and Peter Unger, notable exponents of this view, might be labelled ‘extremists’. They argue that because of the state of the world, and the sheer amount of good we could do, that we are morally required to revise our behaviours to do much more good.

Before looking to the extremists, I first examine an argument by Sobel. He also claims that the objection fails, i.e. that we should not reject a theory (consequentialism, in his case) on the grounds that it is too demanding. Unlike the extremists, rather than making arguments that morality *is* particularly demanding, Sobel rejects the objection on methodological grounds.

I then examine Singer’s position, before rejecting it and examining Unger, who has a more rigorous methodology to argue for similar conclusions. Finally, I give reasons why not only Unger is mistaken in his extremism, but also general reasons why alternative methods to dismiss demandingness objections should be disregarded.

## 2.1 Sobel and the ‘Impotence’ of Demandingness Objections

David Sobel provides an argument suggesting that demandingness objections are ‘impotent’, and that moral theories should either be rejected on different grounds or not rejected at all. In doing so, however, I suggest that he makes use of a non-standard use of ‘demandingness’, which is problematic for his argument, and that the notion of ‘demands’ are of more moral significance than he suspects.

It must be understood that his discussion only relates to maximising act

consequentialism.<sup>2</sup> He claims that this objection is useless, as it must “pre-suppose the truth of prior and independent breaks with Consequentialism” (2007, p.1). Sobel notes that the structure of the demandingness objection is obscure, as are explanations as to why it is supposed to have force. Once the rationale for the Objection is exposed as faulty, Sobel suggests that any version of consequentialism weakened to accommodate the demandingness should also be undermined. He proposes to do so by providing the structure of demandingness intuitions (2007, p.2).

Sobel claims that understanding the “structure of such intuitions” will be sufficient to demonstrate the impotence of the objection (2007, p.2). He sees costs to well-being as the most important sacrifice required by consequentialism, though Sobel admits that the costs<sup>3</sup> of consequentialism can also terms of “life projects” or that it “narrows one’s options significantly”. Sobel also restricts the objection such that the cost of not holding “deontological rights or libertarian property rights” cannot count.

Sobel provides an understanding of the general form of the Objection:

“Morality, properly understood, should not take over our lives, at least in circumstances such as we face these days, but should be compatible with a range of attractive and self-directed lives, including lives that involve a serious commitment to family, friends, or non-moral projects.”

(2007, p.3)

Consequentialism, it is then noted (again, this is strictly *maximising act consequentialism*) does not permit one to live a life of this sort, thus should be rejected. Sobel then provides us with an example intended to demonstrate a problem with this sort of objection. Sally needs one of Joe’s kidneys to live.<sup>4</sup> Joe has two healthy kidneys, and consequentialism claims

---

<sup>2</sup>Sobel does see that demandingness objections serve as a motivation for retreating from this position though (2007, p.1), so problems for the objection here may undermine some rationale for adopting a less demanding theory, such as Scheffler’s move to agent-centred prerogatives (1982, *passim*) or Hooker’s move to rule-consequentialism (2000).

<sup>3</sup>It is notable that even extremely early in his article, when discussing what consequentialism asks of people, Sobel puts this in terms of *costs* which is the position I shall later be rejecting.

<sup>4</sup>Sobel seems to overestimate the damaging effects of donating a kidney – he assumes that it will result in a significant reduction in the quality of Joe’s life, which does not seem to be borne out by empirical evidence, as is pointed out by Larissa MacFarquhar in

he should give one to Sally. However, demandingness intuitions suggest, the “size of the cost makes the purported moral demand that Joe give the kidney unreasonable, or at least not genuinely morally obligatory on Joe” (2007, p.3). But then, Sobel suggests, Sally may complain “about the size of the cost that a non-Consequentialist moral theory permits to befall her”.

Sobel then asks “Why is Consequentialism too demanding on the person who would suffer significant costs if he was to aid others as Consequentialism requires, but non-Consequentialist morality is not similarly too demanding on Sally, the person who would suffer more significant costs if she were not aided as the alternative to Consequentialism permits?” The “obvious answer” offered by Sobel is that we accept that “the costs of what a moral theory requires are more demanding than the costs a moral theory permits to befall the unaided, size of cost held constant” (2007, p.3). Sobel seems to understand the demandingness objection simply as an appeal to costs, or, more specifically, of costs to the moral agents. This raises the question of why the costs to moral agents are of more importance than costs to moral patients.

Sobel’s argument seems to be something like this:

1. The Demandingness Objection requires an acceptance of the premise **costs required by a moral theory are more demanding than costs permitted** (2007, p.3).
2. The Objection does not justify this premise (2007, p.4).
3. If this premise is not justified, it is presupposed, which begs the question against Consequentialism.
4. Therefore, the Demandingness Objection begs the question against Consequentialism.

In responding to this, we might note that in an ordinary – or standard – use of the word “demanding”, the premise is obviously true, thus not in need of justification. To establish this, let us think about what makes something demanding. The types of things that are demanding are those that issue demands. Obligations can be demanding. A situation can be demanding

---

her discussion of altruistic kidney donations (2015, p.193) – but we can assume that cases with relevant features do exist.

if it involves demands being issued. Jobs can be demanding if they make many or stringent demands upon the employees performing them.

All demands also carry with them an assumption of authority; they are more than mere requests. A boss can make demands of an employee, as they are in a position where it is legitimate to require certain behaviours. A government official may send a letter demanding that tax is paid, as the government has the authority (or at least purports to) to legitimately require this. Schedules can be demanding because they make demands on one's time, but only when the source of the schedule has some authority. With demands, there is a sense that the act in question *must* be performed, according to the source of the demand. Despite the normative authority, or claimed normative authority, of a demand, an agent may actually reject, ignore or otherwise fail to comply with a demand. Though it may be bad for your career, you can reject your boss's demands. The same is true with demands of morality. Maximising act utilitarians will believe that many of us *do* ignore the demands placed upon us, when we do not give as much as we could to effective charities. The existence of a choice, whether or not to acquiesce to these demands seems crucial.

There are, of course, looser uses of the word, by which we may say that a condition/situation is demanding, in situations where the above does not apply. We might say "life can be really demanding sometimes", just to mean that it can be difficult.<sup>5</sup> Perhaps we may say that a person is really demanding when they merely make requests, or we could say some work of literature is "a demanding book to read".<sup>6</sup> These would be non-literal uses of the term though, if they do not refer to actual demands and the experience of choosing to perform them. Similarly, if one is not able to ignore/reject/flout the performing of a certain action, then it cannot be demanding.

Now we may return to costs permitted and costs required. When a cost is permitted by morality, morality makes no demand that it be prevented. For example, in a healthcare case, where someone *could* save the life of another, but at some extreme cost (sacrificing a limb or an organ) but this is not required, no one is subjected to a demand. When a cost is *required*

---

<sup>5</sup>This point is also made by McElwee (2017, p.99).

<sup>6</sup>Singer and de Lazari-Radek make this claim about Sidgwick's *Methods of Ethics* to explain why it is less widely read than Mill's *Utilitarianism* (2014, p.viii).

by morality, on the other hand, a demand is made, impelling someone to *do* something. It is thus obviously true that costs required are more demanding than costs permitted, as costs required do make demands, whereas those permitted do not.

Sobel can clearly respond to this by asking what is so special about demands. He does, in fact, note that this ‘Linguistic argument’ can be made, but is useless without any vindication of “the special importance of the costs that a moral theory requires rather than what it allows” (2007, p.8). This is what his rejection of demandingness objections comes down to. Without some independent reason for thinking that demands are of moral import, the objection can not get off the ground. And, he suggests, the objection makes no claims about why demands are important. So, the objection either needs additional support, which would have to come from premises incompatible with consequentialism, *or* it simply begs the question.

However, it seems that Sobel simply neglects the intuitive pull that there *is* something special about demands in this sense. That this is the case has a great deal of plausibility. What his challenge seems to require is some support for the claim that there is something special about demands – costs required of people as *agents* – that can justify their having a different role in moral thinking than costs incurred to patients. This is a peculiar way of thinking about this. Only a consequentialist (and only one of a certain sort) would *not* think there is something special about agents. The practice of making moral decisions, determining what the moral features of an individual’s situation are and how they bear on the moral landscape – this occurs in this minds of *agents*. Similarly, we regard moral *motivation* as an important fixture in our thoughts about morality, but this is only (directly) relevant to agents. Our reactive attitudes towards people also depend on this distinction. Praise or blame are properly felt towards agents, as the persons able to affect the situation, but not towards those who are mere patients.

McElwee responds to Sobel’s challenge to identify something special about costs a demand imposes upon agents primarily in terms of reactive attitudes. He claims that our belief that an action is wrong involves, in part, a judgment that “it is appropriate to feel blame towards” them, and once this is appreciated our judgments about obligation must “take into account

standard patterns of human motivation” (2017, p.99).<sup>78</sup>

It is the consequentialist who takes an unusual view of these factors, and must explain why these seemingly important asymmetries are subordinate to a hedonic calculus.<sup>9</sup> Of course, Sobel is right that supposing that demandingness is important for limiting (or forming) obligations requires a break with consequentialism. But *that’s the point*: we *do* think demands are of a distinctive moral kind!

Of course, we can offer further explanations for why these asymmetries are important. Alternatively, we can see the overly demanding verdicts as a symptom of some other faulty feature of a moral theory. Liam Murphy, for instance, locates the *real* problem in unfair imposition of responsibility (2000, p.99).<sup>10</sup> Attempts like this to dissolve the objection, however, still start with the notion that we have a problem; our given moral theory is too demanding. This itself is enough for an objection of this kind to be made. That we might suppose the problem is caused by failures of different theoretical apparatus does not mean there is not a clear problem we are alerted to by overdemandingness.

Given the asymmetries noted, and the clear intuitive force of demandingness objections, the onus is on people like Sobel to demonstrate why there *isn’t* something special about demands. Lacking such an explanation, the complaint against various moral theories or principles that they are too demanding retains its force.

Now I turn to a paradigmatic extremist: Peter Singer.

## 2.2 Singer

### 2.2.1 Singer’s Arguments for Strong Demandingness

Peter Singer has advocated the adoption of a ‘principle of sacrifice’. Though this principle seemed extremely *prima facie* plausible, after some evaluation,

---

<sup>7</sup>The conceptual link between obligation and blameworthiness is explored further in chapter 5.

<sup>8</sup>Woollard also responds directly to Sobel’s challenge, and focusses more on the “distinctive structure” of moral demands in demonstrating what makes them special. Given that an agent must *choose* to either make some sort of sacrifice or fail morally, it should be “reasonable to expect an agent to conform to morality” (2016a).

<sup>9</sup>John Harris makes an interesting attempt at this type of undertaking in “The Survival Lottery” (1975).

<sup>10</sup>Murphy’s response to the problem is discussed in section 4.2.

it becomes apparent that acceptance of this principle would require all of us who do not live close to the poverty line to radically alter our lifestyles. Singer's argument has spawned a vast amount of literature, revitalising the question of how demanding a moral theory should be.

Singer begins "Famine, Affluence and Morality" by reminding us of the vast quantities of easily-relievable suffering in the world, focussing on the famine affecting Bengal at the time he was writing. With this in mind, Singer provides us with a principle that he supposes we would all intuitively accept:

**The Principle of Sacrifice:** <sup>11</sup> If it is in our power to prevent something very bad from happening, without thereby sacrificing anything of comparable moral importance, we ought, morally, to do it.

(Singer, 1972, p.231)

This principle does have the virtue of sounding extremely plausible. It becomes clear, however, after some thought that the principle inevitably (in any world with such rife amounts of easily preventable suffering, like ours) will entail that we all ought, morally, to devote enormous levels of our time/money/energies to preventing suffering. This is easily seen when considering any choice between everyday activities. When choosing what to do for dinner, perhaps I think it'd be nice to try the new restaurant, but then I realise that I could buy some pasta and cook it myself for much less money, and donate the rest to a charity, which could use the difference to prevent a child going blind. It seems obvious that my sacrificing trying the new restaurant in favour of a cheap meal at home is not morally significant. It is also difficult to deny that a child becoming blind is "something very bad". This procedure can then be applied to any purchase one makes. Whenever a premium good can be replaced by a cheap alternative, or dispensed of altogether, the *Sacrifice* holds that we ought to do so.

Thus with just two additional premises,<sup>12</sup> *Sacrifice* entails a very demanding conclusion. The first additional premise concerns the badness of certain types of suffering. It simply claims that the suffering of those in the developing world from starvation or easily curable diseases falls into the *very*

---

<sup>11</sup>Henceforth *Sacrifice*.

<sup>12</sup>Richard Miller makes the point that these implicit premises are needed to arrive at the demanding conclusion 2004, p.357.

*bad* category. I take this as obviously true. The second additional premise states that we can prevent this very bad situation by doing something of a non-comparable moral importance – namely, donating to various effective aid agencies.

At previous points in history, this could have been questioned. For instance, in his response to claims that utilitarianism is too demanding, Sidgwick argued that one could favour actions that increased one’s own happiness disproportionately to others because of our knowledge of our own happiness. He claimed this because of our privileged practical and epistemic positions with regards to our own happiness, that is, “each man is better able to provide for his own happiness than for that of other persons, from his more intimate knowledge of his own desires and needs, and his greater opportunities of gratifying them (1907, p.431).<sup>13</sup>

Richard Price made a similar point in his *Discourse on the Love of Our Country*, claiming that even though we should adopt a view of universal benevolence, we naturally focus on ourselves and our friends/families, as “we can do little for the interest of mankind at large” (1789, pp.10-12).

However, in the world as it is today, with charities that work effectively to prevent diseases, provide healthy water etc. we clearly can prevent a great deal of suffering while sacrificing comparatively little. With our greater knowledge of the extent of global poverty, and the availability of institutions that can effectively relieve it, we can not able to deny this additional premise.<sup>14</sup>

The *radical conclusion* we are left with is that affluent people – most people in the western world – should be donating as much as they possibly can without impoverishing themselves or a dependent, that “everyone has a duty not to spend money on luxuries or frills, and to use the savings due to abstinence to help those in dire need” (Miller, 2004, p.358).

As the additional premises are so clearly true, we must either accept this radical conclusion, or deny *Sacrifice*<sup>15</sup>. Most have opted for the latter

---

<sup>13</sup>Singer and de Lazari-Radek note that this response no longer seems tenable in their discussion of Sidgwick (2014, p.323).

<sup>14</sup>Both of these additional premises – regarding the badness of the suffering and our ability to prevent it without sacrificing anything of comparable importance – are made explicit and expressed in a formal version of the argument in Singer’s *The Life You Can Save* (2010, p.15).

<sup>15</sup>Even people like Dorsey, if they do not deny the principle, are committed to the conclusion that we have a *moral* duty to act this way, even if they deny that we have

option, but Singer attempts to justify this principle.

Other than its prima facie plausibility, Singer offers little to defend his principle. He does, however, illustrate how his principle works by making use of his famous ‘drowning child’ example:

“[I]f I am walking past a shallow pond and see a child drowning in it, I ought to wade in and pull the child out. This will mean getting my clothes muddy, but this is insignificant, while the death of the child would presumably be a very bad thing.”

(1972, p.231)

Singer takes our acceptance of the moral verdict in this example for granted before defending against to potential disanalogies which might arise between the drowning child case and the case of our obligations to the global poor. When the implications of the principle with regard to these disanalogies are appreciated, the resulting moral obligations become extremely burdensome.

The first disanalogy concerns distance. In the pond case, the agent in question is very close to the event, whereas in the case of world hunger, this normally occurs considerably further away from affluent readers of Singer’s argument. Singer suggests that there is no morally salient reason why proximity should make any difference. Assuming we are able to help the would-be-victim just as easily, it should not matter. Because of our commitments to “impartiality, universalizability, equality, or whatever, we cannot discriminate against someone merely because he is far away . . .” (1972, p.232), as surely that does not affect their moral status. Due today’s world being a ‘global village’, due to efficient transport links and the ease and speed of information-transmission, we are able to help those suffering from a great distance.

The second potential disanalogy Singer raises concerns the number of potential ‘helpers’ there are. In the original pond case there is only one moral agent who sees the drowning child, whereas with regards to the global poor, there are many in a similar position equally able to help. Here Singer seems to *defend* the principle with a modified ‘drowning child’ example.

---

an *all-things-considered* duty. Denying overridingness might, however, have the effect of making this more palatable, as we shall later see.

Simply adding additional bystanders, Singer notes, does not make it any more permissible for the onlooker to simply watch the child drown.

Having defended his principle against rejection for those reasons, Singer moves on to suggesting a change in the way we see giving. Whereas donations to famine relief seem like optional good things to do, perhaps of a supererogatory sort, Singer claims they are in fact duties (1972, p.235). This conclusion does seem unavoidable if we accept Singer's principle.

### 2.2.2 Rejecting Singer's Argument

I think there are many good reasons to reject Singer's principle. Several of these apply not only to Singer's argument, but also to the general extremist position. I engage with some of those arguments in section 2.4. In this section I respond to issues specifically regarding Singer's argument.

We might first note that Singer's argument is a fairly strange one. He presents us with *Sacrifice*, the main supporting feature of which is its "uncontroversial appearance" (1972, p.231). He then provides us the example of the drowning child, for which the principle seems to let us arrive at the correct verdict. It seems as though we are expected to make the inference that if we think we have an obligation to save the child, then we should accept *Sacrifice*. Leif Wenar actually accepts this inference, claiming that those "who believe they ought to save the child should endorse the moral principle" (2010, p.105). But it is obvious that, without any additional support for the principle, this is a bad inference. There are many principles which do not entail *Sacrifice* that would arrive at the verdict that we ought to save the child,<sup>16</sup> and Singer provides no reasons to favour *Sacrifice* as opposed to those alternatives.

This hardly seems like the rigour we might expect or desire when considering a principle which, if accepted, would completely change the way we view our moral landscape.<sup>17</sup> We might, for example, wonder if we can provide cases where our intuitions go against what *Sacrifice* would suggest. If the principle is rejected, the argument fails.

One feature of Singer's principle we might consider problematic is that

---

<sup>16</sup>Miller's *Sympathy* principle (2004; 2010), which is discussed in section 4.1, and Murphy's Collective Principle of Beneficence (1993; 2000), discussed in 4.2, are examples of such principles.

<sup>17</sup>This point is also made by Travis Timmerman (2015, p.207).

it only pertains to *actual* sacrifice. Often, the morally comparable features will not *obviously* be sacrificed by some conduct, but might well be *risked*.<sup>18</sup> Obviously there will be a difficulty for the agent to know whether they are *sacrificing* or merely *risking* sacrificing whatever good is in question. This is clearly a very important consideration. Consider a modified version of Singer’s pond case. If the onlooker realises she is a fairly poor swimmer, and thinks there is a non-negligible chance she too would drown<sup>19</sup> this is likely to affect our assessment of her obligations. Even if she actually would have successfully saved the child, the existence of risk definitely has the ability to alter the verdict.

This point is an important one too. We could consider a weakened principle:

***Sacrifice\**** If it is in our power to prevent something very bad from happening, without thereby significantly<sup>20</sup> risking sacrificing anything of comparable moral importance, we ought, morally, to do it.

If this principle were adopted, however, it would not require — as Singer suspects that the original principle might — agents to reduce their welfare up to “the point of marginal utility” (1972, p.234). Even if the bar for what might count as significant is held fairly high, it seems that leaving oneself with no savings would — for many at least — impose significant risks of destitution, which would presumably be of comparable moral significance.

This modification to Singer’s argument, while definitely making it less demanding, would still (as long as ‘significantly’ is taken to do enough work to make *Sacrifice\** remotely plausible) leave an extremely demanding account of the moral obligations of affluent people.

If it seems unfair to characterise Singer’s argument as captured entirely by this paper (which, admittedly, is over 40 years old), we should remem-

---

<sup>18</sup>This is implicitly recognised as problematic by Richard Miller, whose own principle — a more moderate alternative to Singer’s proposal — includes the caveat that one’s duties of beneficence should not “impose a significant risk of worsening one’s life” (2004, p.359).

<sup>19</sup>Exactly what this chance would be is difficult to say, but I will take it for granted that at *some degree of risk* we accept that she is not required to try to rescue the child — and perhaps even that she *should not* try to rescue the child.

<sup>20</sup>Following Miller, I have used *significant* risk, rather than any risk at all. Given distant possibilities, *any* risk would make the principle extremely weak. This does come at the cost of some vagueness, as it is a difficult matter to say what might count as *significant* risk.

ber just how crucial the principle is for Singer in deriving the demanding obligations (and our re-evaluation of what is required by duty, rather than charity). In recent works, Singer has still defended this principle (e.g. 2007, p.476, 2010, p.15, 2015, p.28), so the argument is at least one that he considers to have withstood extensive criticism.

I will now attempt to provide cases that could serve as counterexamples to Singer's principle. A good counterexample to Singer's principle must include a situation wherein "it is in our power to prevent something very bad from happening, without thereby sacrificing anything of comparable moral importance" but we do not think we are morally obligated to do so. Singer's opponents could cite the obligation to give substantial aid to the world's poor as a counter-example, but this would certainly be question-begging, and various other types of example seem possible. I provide four types of cases that are at least good candidate counter-examples to *Sacrifice*:

#### **Example type one: Obvious division of labour**

Consider a small village, before mobile phones became commonplace. The residents are all reasonably well-off - not rich, but all with some savings - and one, Jennifer, is about to leave for a big trip she has saved up for. She is very excited and has her travel booked. She sets out, needing to leave immediately the morning of her departure, and when she reaches the village gate, she runs into Frederick, a man from her village who appears dismayed. Apparently Frederick has just been diagnosed with a horrible disease. Unless he can afford an operation by the end of the week, he will die. The operation costs 10,000 Crowns and will guarantee his recovery. Jennifer knows that she actually has 10,000 Crowns on her person, but it is everything she has saved for her trip. She also knows that the 100 people in the village can all afford to give Frederick some money, probably at least a 100 Crowns each, though she is not *sure* that they would.

If Jennifer gives Frederick all her money, she will not be able to go on her trip, which she has looked forward to for over a year. However doing so, it seems, would prevent something bad from happening (Frederick dying). If Jennifer just gives Frederick 100 Crowns (or maybe 500, just in case a few people in the village are a little short), she won't have *prevented* the bad

thing happening. We know she hasn't prevented it from happening, as it could still happen; everyone in the village might uncharacteristically behave horribly.

Jennifer would readily acknowledge that her trip is not of comparable moral importance to Frederick's life. Despite this, it doesn't seem like she does anything wrong if she merely helps out with a reasonable share of the cost for Frederick's treatment. It might be nice if she did more, but it certainly wouldn't be her duty.

It seems that our intuitions about this case will be different to the pond case because the cost is so easily divisible. The case is designed such that Jennifer is leaving the village (and in a considerable rush) in order to remove the option of conferring with her fellow villagers and ensuring compliance with a fair or adequate donation by everyone. Unfortunately for Singer, this does seem to be analogous to the situation we find ourselves in with regard to world poverty and any given individual in the developing world who might be saved in the near future by an influx of money to an effective charity. If the amount of money needed to prevent a disease is calculated, it could be divided (using some fair distribution procedure) among everyone in the affluent world.

Singer might reply, however, that this example is not a reasonable parallel to the issue of world poverty for two reasons. Firstly, ensuring compliance over the western world seems impossible (without governmental intervention, at least), whereas in the friendly close-knit village, it may be expected. Secondly, the quantity required might be such that large contributions would be required from everyone anyway, which may not be suitably similar to Jennifer's reduction in travel money. As a consequence, though *Sacrifice* seems to fail here, a similar principle generating Singer's desired obligations in the case of world poverty might better account for our intuitions here. For example, *Sacrifice* could be supplemented by a clause stipulating that one should prevent the bad thing from happening, unless one can partake in partial compliance with a broader solution (e.g. everyone paying a fair amount) and expect full compliance from the wider community. A clause like this might be cumbersome, but maintain the spirit of Singer's principle, and similarly generate the obligations of *Sacrifice*, because, as he rightly points out, we do not expect full compliance in our current situation (1972,

p.233).<sup>21</sup>

### **Example type two: Clear iterations to big obligations**

Consider the following situation: You are informed by a reliable doctor that a friend of yours will die immediately without your attention. Your friend is young, well-liked and enjoys his life. His dying right now would clearly be a very bad thing, and you can prevent it. Right now, you are told, all you need to do is go to his side and pay very minimal attention to him to ensure that he will not die now. Your having to be by his side and pay minimal attention is surely not of comparable moral importance to his dying right now, so it seems clear that you ought to help.

However, after helping once, you discover that your friend will be in this position for the rest of his life. For some reason, which the doctor explains but you don't understand, only you can help. Your friend needs your constant presence (or presence say, every ten minutes) to remain alive. He cannot be moved from the hospital, but he still really enjoys his life there, as his favourite activities (e.g. computer games, or chess perhaps) are all available to him there. For every time you must make your decision, however, your decision procedure, if governed by *Sacrifice*, will require you to attend to your friend. Every time you make the decision, your friend dying then would (or at least we may suppose) be a very bad thing.

*Sacrifice* seems to require that, through the extended iterations, you are morally obligated to attend to your friend for the rest of his life (or the rest of yours). This seems like a wildly implausible conclusion. Though each individual act is not a big sacrifice – it only asks that you spend an additional ten minutes there – the commitment required by the cumulative sacrifices is enormous. We may be reminded of Judith Jarvis Thomson's discussion of her violinist. If you wake up strapped to a violinist, who requires the use of your body to survive for the next nine months, Thomson claims that you are not obliged to help (though it might be nice to do so). If the example is

---

<sup>21</sup>I will not return to considering an option like this, because an amendment like this will still leave the principle vulnerable to all the other example types given.

modified however, so that the violinist only requires you to stay in bed with him for one hour, “you ought to allow him” to do so (1971, p.59). Cases such as this highlight the difficulty in determining how much one ought to do. Unfortunately for Singer’s position, iterative concerns like this do apply to the issue of our obligations to the world’s poor, as it is such iterations which lead to the radical conclusion.

While this example is equivalent in form to the general iteration worries mentioned previously,<sup>22</sup> I hope this particular example is useful in a number of ways. Firstly, it seems just as obvious that the conclusion is mistaken as it does in the case of the radical conclusion. It also should (I hope) be at least as compelling that you should help in the first instance, as it is that an agent should refrain from a luxury good in order to send money to famine relief charities. If I am correct about the strengths of these intuitions, we should have even more reason to be suspicious of the iteration of obligations. If we have good reason to suppose an obligation arises in one case, but does not apply in every following case despite the relevant reasons being exactly the same, we should be very suspicious of a principle which is consistent with obvious iterations to big obligations.

I do suspect that the best resolution to this problem is to adequately define “something very bad” such that the death of an innocent, young, happy person will not necessarily suffice. In this example, the friend dying *one hour earlier* than he would otherwise, for example, might not count as very bad. While I think this problem might be avoided by properly defining terms, this was not really even attempted in Singer’s paper.<sup>23</sup> This should at least make us reluctant to accept the argument without a little further pressing.

### **Example type three: Presumption<sup>24</sup>**

You discover that an acquaintance of yours has fallen, through

---

<sup>22</sup>e.g. section 1.2.1

<sup>23</sup>In fact, the brief description of what might count as *bad* do not help in this situation. He suggests that “suffering and death from lack of food, shelter, and medical care are bad” (1972, p.231). Though I did not stipulate that the death would involve suffering, we can amend the example to include as much suffering as would be required to meet that criterion.

<sup>24</sup>For this type of example, I am indebted to Rowland Stout, who argued (at a conference on supererogation in Dublin in June 2014) that presumption differs from supererogation, not by the type of act committed, but due to the relationship one has with the recipient.

no fault of their own, under hard times. They have borrowed money from a violent loan shark, who may break their legs if they don't pay up. You do not know the person particularly well, and arrived at this knowledge through a series of unusual coincidences. Perhaps you volunteer for a confidential, listening and information service for people undergoing stressful situations, but you have inadvertently discovered who the caller is. You also know that the person in question has a wealthy, forgiving and supportive family, who they could (and you think should) ask for assistance.

Though you *could* help the person with very little difficulty to yourself, it would be presumptuous for you to do so, as you do not know them particularly well. Also, though you know that you could categorically prevent the bad thing (their legs being broken), the optimal situation (or so you have surmised and have a very strong conviction of) is that they return to their family and seek their support, financially and emotionally, as this will not only prevent one bad event, but perhaps future issues as well.

I have supposed that you would not be under an obligation to help in this situation because it's *someone else's* job to do so. It would be presumptuous of you to help. It could be argued that even if the family could not help out, it would be supererogatory of you to help rather than obligatory. As it is a stranger and because one lacks the special relationship (which you know he bears with others with the ability to help), this may affect the duty (or lack thereof) one has towards him.

If it is accepted that there are cases wherein one could easily prevent a very bad thing from happening, but where it would be inappropriate (or wrong even) to do so because it would be overly presumptuous, *Sacrifice* must be rejected (or at least modified).

#### **Example type four: Desert**

Another acquaintance of yours has also become tangled with loan sharks several times in the past. On several previous occasions, you have bailed him out, despite the fact that he has never been a particularly close friend. You have done so sheerly because you knew him and knew that he did not have family or friends who

could obviously help. The last time this happened, you warned him not to become involved with such nefarious knaves again and that if he did you would not help.

While on vacation, several hours away from this acquaintance, you get a phone call. This acquaintance is in trouble again, and if he doesn't get the money, in cash, by the end of the night, the loan sharks have threatened to break his legs. You could make it back in time to help, providing him with the cash, which wouldn't substantially hurt your wallet or take a colossal amount of your time.

The fact that this agent has brought the situation upon himself – he is extremely culpable – seems to diminish anyone else's supposed duty to help. While it might be very bad if his legs are broken (some might argue - perhaps somewhat callously - that this would be justice, rather than a bad thing), you certainly don't seem under an obligation to help.

The response to this case that it would not be a bad thing for this person to have his legs broken might be made by some. However, it might be difficult to reconcile with other positions (e.g. utilitarianism). The acceptance of *Sacrifice* with a specialised notion of what 'bad' is might thus require additional theoretic commitments (or at least rule some out).

In all of the above examples it might be thought that to help would be good, but not required of the agent in question. This relates to the third objection, questioning why the sacrifice would be obligatory, rather than supererogatory. Hopefully, the specific features of the example explain what about the particular example makes it the case that one is not duty-bound to help.

I do think these counterexamples cast serious doubt upon *Sacrifice*. It is certainly the case that there exist similar variant principles that would generate extremely demanding obligations to the world's poor while also harmonising with our intuitions in these types of example, but this need not be discussed here. This will be worth bearing in mind after chapter 3,<sup>25</sup> the conclusions of which will, if accurate, also provide a motivation to dismiss *Sacrifice*. It is enough here to reject *Sacrifice* as a general principle of beneficence. I shall suppose that *Sacrifice* fails to be a satisfactory moral

---

<sup>25</sup>These are elaborated further in the final chapter.

principle. Due to the central role this principle plays in Singer's generation of the huge demands, this suggests we should reject Singer's argument. As I consider later,<sup>26</sup> there are also considerable independent motivations for rejecting extremely demanding conclusions.

## 2.3 Unger

Singer attempts to convince us of his position via one apparently obvious principle. Accepting this, he draws conclusions which then seem counterintuitive. Unger has a similar project, but attempts to demonstrate certain features of our psychology which he hopes that we all will accept lead us to make poor moral judgements. Once we notice these, like Singer, he suggests that we will see that the demands of morality are significantly more extensive than commonly appreciated.

### 2.3.1 Peter Unger's Argument for Demandingness

Unger begins by noting what he takes to be a puzzle. This is brought out in two cases and our intuitive judgements about those cases.

The first case Unger provides is **Vintage Sedan**:

“Not truly rich, your one luxury in life is a vintage Mercedes sedan that, with much time, attention and money, you've restored to mint condition. In particular, you're pleased by the auto's fine leather seating. One day, you stop at the intersection of two small country roads, both lightly travelled. Hearing a voice screaming for help, you get out and see a man who's wounded and covered with a lot of his blood. Assuring you that his wound's confined to one of his legs, the man also informs you that he was a medical student for two full years. And, despite his expulsion for cheating on his second year final exams, which explains his indigent status since, he's knowledgeably tied his shirt near the wound so as to stop the flow. So, there's no urgent danger of losing his life, you're informed, but there's great danger of losing his limb. This can be prevented, however, if you drive him to a rural hospital fifty miles away... Now, if you'd

---

<sup>26</sup>In section 2.4

aid [him], you must lay him across your fine back seat. But, then your fine upholstery will be soaked through with blood, and restoring the car will cost over five thousand dollars. So, you drive away. Picked up the next day by another driver, he survives but loses the wounded leg.”

(1996, pp.25-26)

Unger notes that most people, when confronted with this case, share the intuition that this behaviour is despicable. Our judgements of someone who acts in this way are not likely to be favourable.

This is then compared with a second case: **The Envelope**:

“In your mailbox, there’s something from (the U.S. Committee for) UNICEF. After reading it through, you correctly believe that, unless you soon send in a check for \$100, then, instead of each living many more years, over thirty more children will die soon. But, you throw the material in your trash basket, including the convenient return envelope provided, you send nothing, and, instead of living many years, over thirty more children die soon than would have had you sent in the requested \$100.”

(1996, p.25)<sup>27</sup>

In contrast to the first case, we do not find this case to promote a hostile reaction. Despite many more people (thirty, rather than one) suffering a much greater loss (death, rather than loss of a limb) which could be prevented by a much lower cost (\$100 rather than \$5,000), the typical intuition in this case is that the agent in this case has done nothing wrong. We feel nothing like the disapprobation we feel in the first case. Therein lies the puzzle. Why would our intuitions judge the conduct in first instance so harshly, but the second so leniently?

---

<sup>27</sup>We might note that the particular figures for how much it *actually* costs to save a life are estimated to be significantly higher than this. In *The Life You Can Save*, Singer estimated that we could confidently judge that the best charities could save a life for “between \$200 and \$2000” (2010, p.103). This would, however, not phase Unger or affect his conclusion, as he maintains that “it’s badly wrong not to provide vital aid even if it costs *many thousands* of dollars to less by just *one* the number of distant children who’ll die young rather than live long” (1996, p.145)

In attempting to show that this puzzle is not easily resolved, Unger examines several factors that potentially ground asymmetries between *Vintage Sedan* and *Envelope*, and amends the cases slightly for each of these factors to demonstrate that our intuitions seem to persist. Among the factors he considers are physical proximity, social proximity, informative directness, experiential impact, the thought of a disastrous further future and whether there are multiple or unique potential saviours. Keeping the cases extremely similar to the original versions of *Envelope* and *Vintage Sedan*, Unger then attempts to remove these factors, (ideally) to show that our intuitions do not change, and thereby that what our intuitions are sensitive to is something different.

Unger eventually claims that the factors that do affect our intuitions in these types of cases are extremely strange, clearly not of moral relevance and because of this, that we should ultimately reject (or ‘liberate ourselves from’) these intuitions. The major factors Unger identifies as the culprits are what he refers to as *projective separating* and *protophysics*.

Projective separating is a term Unger uses to describe the way we group people when we consider certain cases. Typically, and of no moral relevance (Unger claims), we...

“... view the world as comprising just certain *situations*. Likewise, we view a situation as including just *certain people*, all of them then well grouped together within it... viewing all the world’s *other folks* as being only in *other situations*”

(Unger, 1996, p.97)

The upshot of this separation is that when we group one set of people with a problem separately to another, who do not have the problem, we deem it unfair – and morally wrong – to impose losses on the group without the problem in order to aid the group that does.

As well as explaining why we do not negatively judge the conduct in *Envelope* – because the children are in a different situation – and why we do in *Vintage Sedan* (because we *are* in the same situation as the ex-medical student), Unger notes that this explains our responses in other difficult cases. In Philippa Foot’s original trolley case,<sup>28</sup> for example, where an agent can

---

<sup>28</sup>In that discussion, she actually refers to a “runaway tram”, but the case is otherwise the same (1967).

change the track of a trolley to prevent it killing five, but doing so will divert it towards one person, the typical verdict that she ought to do so can be explained by us grouping all the seven in the same situation. Foot notably attempts to explain the intuition in this case by invoking the doctrine of double effect.<sup>29</sup> However, Unger provides a separate example, like the trolley case, but which apparently leads most to a different verdict. The case in question, *Yard*, also features an empty trolley heading towards five people, but where it can be stopped by sending another empty trolley to collide with it. If one does that, the trolleys will both be derailed, fall down a hill and into someone's yard "where they will wreak fatal havoc on the yard's owner, asleep in his hammock" (1996, p.98).

In the original trolley case, we see all the people on the tracks as grouped in the same situation, along with the runaway trolley. The trolley there is a problem for all the people on the tracks. In *Yard*, however, we see the group of people on the tracks as in the same situation as the trolley, but the yard's owner as unconnected.

Unger attempts to demonstrate just how our intuitions are affected by this phenomenon with a several-option trolley case, where a runaway trolley will, without any action from the agent, kill six people. Unger then introduces options at several stages. The options are as follows (and illustrated in the picture on the next page (1996, p.90)):

- A Do nothing, and a trolley hits and kills six people.
- B Press a switch, diverting the trolley to the a track with three people on instead. If you press this, the six people will survive, but the three will die.
- C Press an alternative switch, which will push another trolley, which would divert a second trolley, with two people aboard, to hit the first trolley, stopping it but killing the two people aboard this trolley. This will save six people trapped on the track, but the two people on the trolley die.
- D A fat person nearby is roller skating, with special roller skates you can control with a remote control dial. You can activate the skates and

---

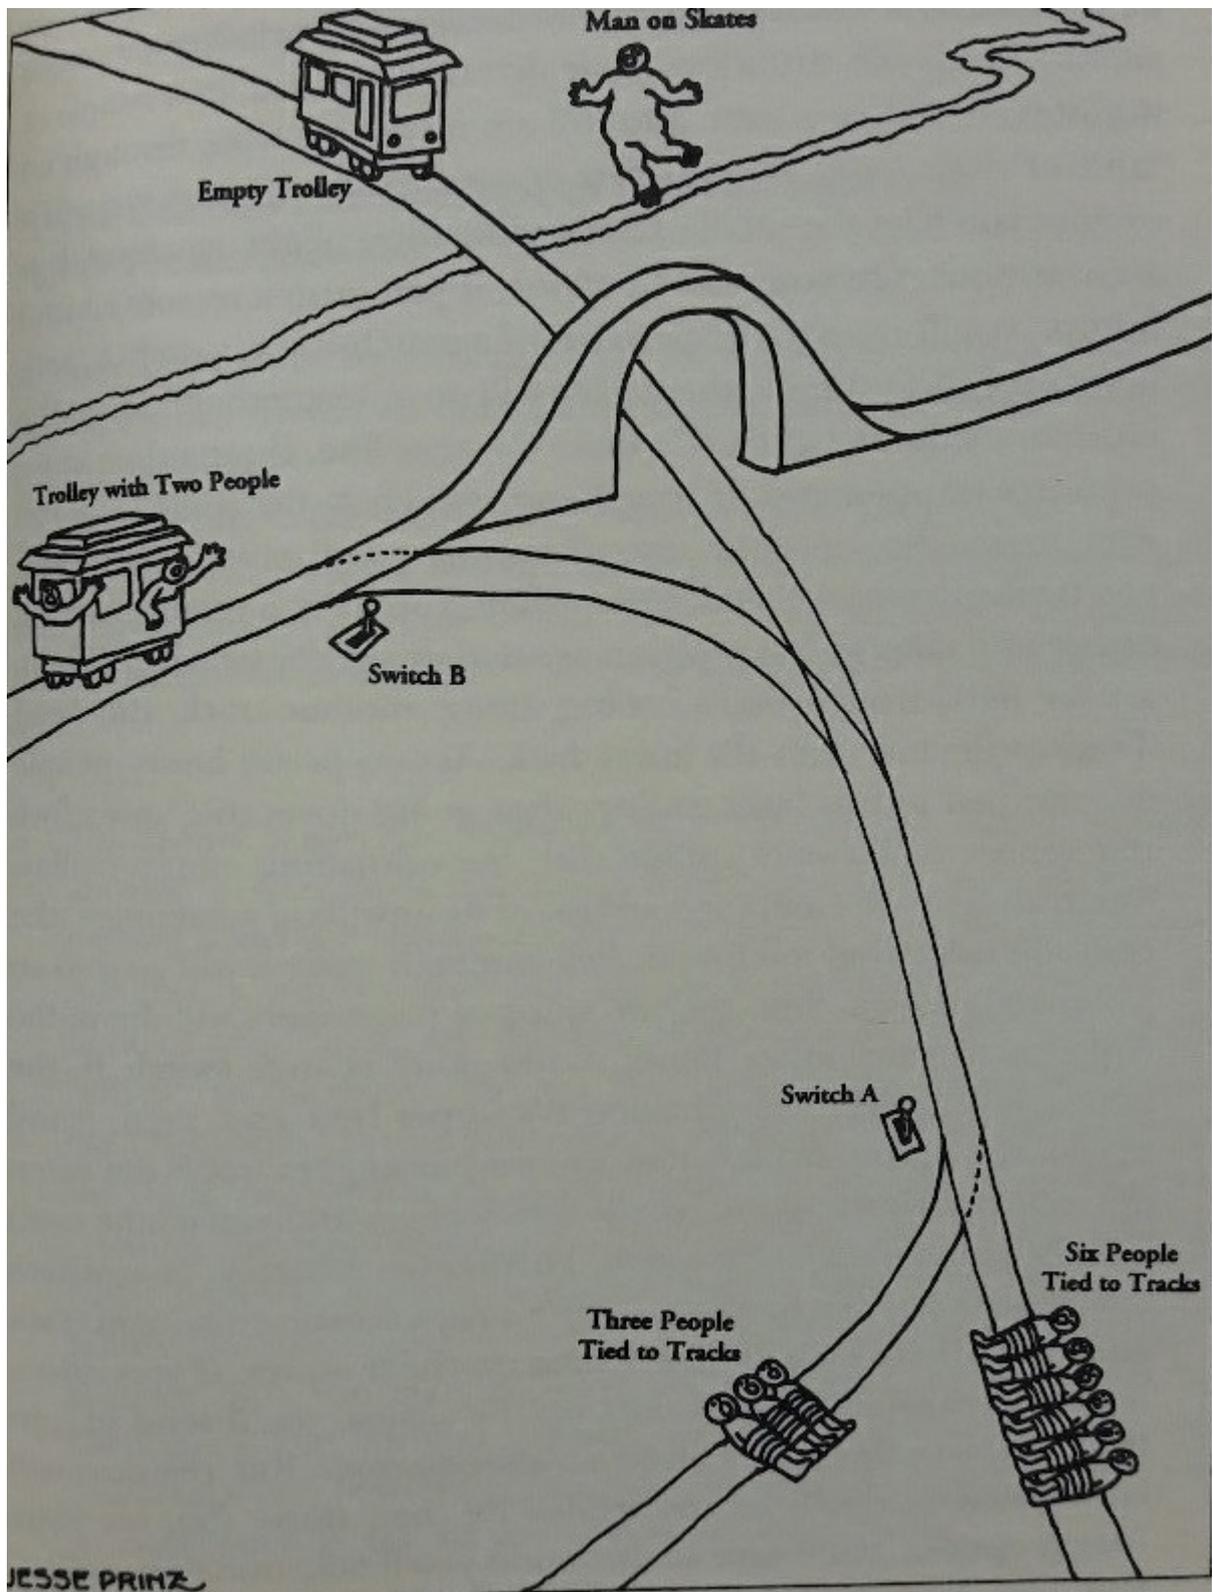
<sup>29</sup>Foot compares the trolley case to one where a judge can frame an innocent person for a certain crime to prevent riots and was engaged in trying to find what makes this impermissible, yet the trolley case permissible.

push the fat person into the trolley. The fat person will die, but the trolley will stop and the six will be saved.

Unger claims that when confronted with *all* of these options, people typically judge that people should choose option D – killing the heavy skater to save the six. That is to say, they have a positive verdict of someone who takes option D. However, if people are only confronted with options A and D, people intuitively think option D is terribly wrong.<sup>30</sup>

---

<sup>30</sup>Unger does note that order effects play a role in this. People have a tendency to attempt to be consistent in the evaluations, so will often show a reluctance to change their verdicts of whether D is wrong given only the addition/removal of options. Despite the differences in verdicts with the differences in presentation, Unger notes the prevalence of the verdicts he describes (1996, p.92).



Projective separation can explain this distinction, because when only confronted with options A and D, the heavy skater appears to us in a totally different situation to the six people on the tracks. When the four situations are presented, however, B seems *close enough* to A, C *close enough* to B, etc. so we do not engage in projective separation. Rather, in that case they are in the same group.

Unger claims that grouping people in our minds in this way is clearly morally unimportant, and that our being influenced by this type of factor is a feature of our thinking we should fight against.

The second psychological factor Unger calls our attention to is *protophysical thinking*. This is the term Unger gives to a variety of unusual ways features of movement affect our judgements. He presents five “protophysical principles” which supposedly distort our intuitions in many cases:

1. When serious loss will result, it’s harder to justify moving a person to, or into, an object than it is to move the object to, or into, the person.
2. When serious loss will result, it’s harder to justify changing the speed of an moving object, or changing its rate of motion, than changing the object’s *direction* of motion.
3. When there’ll be big loss, it’s harder to justify speeding up an object than slowing down an object.
4. It’s a lot harder to justify taking an object at rest and setting it in motion than to justify taking an object in motion and increasing its speed.
5. Should serious loss result, it’s harder to justify imposing a substantial force on an object than it is to justify allowing a force already present (just about) everywhere, like gravitation, to work on the object.

(Unger, 1996, pp.101-2)

Unger again makes use of a variety of examples where our intuitions do not seem to represent our values to illustrate these principles. To briefly illustrate, the first of these principles does some work in explaining our intuitive reluctance to push a fat man (or redirect a fat skater) into a moving trolley, but our acceptance of merely changing the direction of the trolley.

Unger also notes that when just presented with options A and C in the multi-option case, there is a resistance to option C, where one moves *people*, rather than just moving objects (1996, p.102,n.15). Option C does violate the first principle, but because it does not violate any others, Unger suggests that one intermediary case (B), is enough for us to overcome that resistance.

While projective separation and protophysical thinking are the main factors he discusses, Unger notes that there are a wide array of additional psychological features which seem to divert us from making good moral verdicts. Behaviour with typically negative stereotypes (like pushing someone, as opposed to causing them to be moved by remote control), with more psychological distance (where the causal chain between you and the potentially suffering patient is longer) are other such features he mentions (1996, p.105). He also notes that these features often play in tandem, with ‘factors of protophysical thinking [doing] at least some of their deceptive work *through encouraging the work of other* distorting factors, like the factors of projective separating” (1996, p.103).

Unger claims that none of the complex set of strange factors are morally significant. They do however serve to distract us from living morally decent lives. By being alert to these factors and how they might warp our thinking, Unger suggests, we can remedy this and realise that even to live a morally decent life (not even a morally good one) will require much more than is typically accepted.

### 2.3.2 Rejecting Unger’s Argument

Tim Mulgan notes two ways we might resist Unger’s extremely demanding conclusion (2000, p.400). First, we may not share the intuitions Unger claims are typical. Mulgan claims that the intuitions Unger reports do not seem to have been reflected among his students. Kamm also rejects many of the intuitions that Unger sees as revealing the so-called “distortions” in our moral thinking, particularly noting that she does not share the differing judgements in *Switches and Skates* (1999, p.302).<sup>31</sup> Ascertaining what are

---

<sup>31</sup>Mulgan does not go as far as rejecting the intuition himself, instead claiming that he has no firm intuitions in this case. It is notable that many of the cases provided, particularly in Unger’s arrival at the conclusion that projective separation is a significant factor in forming our intuitions, are extremely unusual cases. If we held a view about our moral intuitions which suggested they are less reliable in cases people are unfamiliar with in their everyday lives, this would undermine Unger’s conclusions.

the typical intuitions in this case is clearly an empirical matter, and I will not speculate on this here. In what follows, I assume that Unger's claims about what intuitions people have is generally correct.

Secondly, we might claim that it does not matter even if they are. Mulgan notes that we might respond by suggesting that morality is actually significantly more complicated than Unger suggests – that whether it is permissible to send in the 'heavy skater' may depend on (the seemingly) irrelevant options (2000, p.400). While that option might seem appealing, this would still require a rejection of *some* intuition, namely the second-order intuition that an addition or removal of options we are not going to choose should have no relevance to the permissibility of an action. Unger does note that this second-order intuition is widely shared, so preserving all our intuitions might not be possible.

Another issue with Unger's argument concerns the way he determines that factors are morally irrelevant. When demonstrating, for example, that physical proximity does not change our verdicts, he provides a case very similar to *Vintage Sedan*, but where the distance is greater, and a case like *Envelope*, but the physical distance is greatly reduced. Even if our intuitions do survive these modifications, however, this does not demonstrate that these factors have *no* impact. They may affect our intuitions to a slight degree, such that they only cause a change in our overall verdicts when they are combined with other factors.

This might not be a huge problem for Unger. He notes that our judgments seem to track salience, which is relevant to many of the factors he considers, yet (or so Unger suggests) does not really track our Basic Moral Values (1996, p.29). So long as our intuitions depend upon considerations with no real moral significance, Unger's conclusion remains unthreatened.

Perhaps more seriously, I suggest that Unger's account fails in the following two ways. Firstly, it fails to establish that the 'distortions' really are morally arbitrary. Finally, it sneaks in utilitarian assumptions, which 'common-sense morality' would typically reject.

### **Defending the 'Distortions'**

Unger relies heavily on the claim that projective separation and protophysical thinking (as well as a variety of other unusual psychological phenomena) are morally irrelevant – that they distract us from things that *really* repre-

sent our Basic Moral Values. These features do seem bizarre – how could thinking about people as being members one group rather than another justify killing or letting die in a situation? A response which I find appealing, however, involves defending some of the “distorting” factors that Unger views as responsible for some of our undemanding views. Here we might claim that *even if* Unger is correct about typical judgements, *and* the psychological factors that help explain them, that this still does not seem to be a problem.

Let us consider projective separation. In the comparison between *Trolley* and *Yard*, the typical intuitions Unger describes condone switching a trolley from a track with six people to a track with one person, but condemn sending another empty trolley to collide with it, sending one crashing down a hill to kill a distant yard owner and destroy his yard.

One way we may defend this ‘distortion’ is based on the degree to which we value safety. It seems that people who are on a track – even if it is disused – are in serious danger. The prospect that a trolley could hit them is at least salient, particularly if a trolley is on its way. The yard owner, far away and peacefully napping, has no reason to suspect any danger from this potential source of untimely demise. Because we value the ability to plan our lives, notably to avoid dangers to ourselves and our loved ones, being in a situation which seems safe from certain dangers, and relying on those situations not being changed bizarrely, is important to us. If we can defend cases like this by pointing to something which does seem morally relevant, Unger’s insistence that our intuitions are “distortions” seems, at the very least, less obvious.

Similar justifications can be offered in a typical organ case. Even if one individual’s organs could be used to save five people, we expect that we are safe from being abducted and our organs being harvested. If we accept that imposing a risk upon someone who could have expected to be safe from that risk is morally significant, we go some way to justifying projective separation.

One of Unger’s other cases he uses as a comparison to *Yard* also offers some defence for this type of view, **The Small Missile**

“Six Innocents are trapped on a trolley track. . . [but] trolleys don’t run there anymore. Down the hill from this track and across the road, there’s someone in his own backyard, sleeping in

his hammock. Accidentally, one of the army's very small missiles has been launched. If you *do nothing*, your first option, the missile will land where the six are and, upon impact, it will kill them. . . .By pushing a remote control button, you can reroute this missile, but only in a certain way. Now, if you thus deflect it, then the missile will land in the noted yard and, upon impact, not only will it kill the yard's owner, but it will destroy his entire house and yard. . . .”

(Unger, 1996, pp.98-99)

Unger notes that unlike in *Yard*, most people share the verdict that acting, thereby changing the missile trajectory, thus killing the yard's owner, is a good thing to do. Unger suggests that the reason for the difference is that in *Yard*, we are encouraged by the example to view the napper as separate from the others, but in *Small Missile* all the people are similarly unrelated to the missile.

Unger then suggests that this case diffuses the projective separation, thereby eliminating “morally irrelevant reasons”. However, I have suggested that we can defend this ‘distortion’. While the napper seemed safe from trolleys (particularly compared to the six people on the tracks), he was no safer from this freak occurrence than anyone else.

This way of thinking can also explain why some of the protophysical principles seem to hold sway. Consider the final, and possibly most bizarre-seeming one, that “it's harder to justify imposing a substantial force on an object than it is to justify allowing a force already present (just about) everywhere, like gravitation, to work on the object” (1996, p.102). Forces that are (just about) everywhere are more easily predictable. Agents have the ability to plan for these and alter their lives accordingly.

We might then propose this type of justification – that we take as extremely valuable the ability to predict and plan our lives in typical situations – a little further. We could, for example, accept that burdens which are *very likely* to put into peril agents' abilities to plan and predict the course of their lives, carry additional moral weight, i.e. they require considerable more to be at stake. This might allow us to make the type of manoeuvre that Brad Hooker deems as promising in response to Unger. Hooker suggests that there might be a legitimate moral distinction between costs which are

likely to affect us repeatedly and those which are likely to occur never or extremely rarely (1999), because of the expected effect of complying with demands. Our placing a huge value on the ability to avoid certain dangers – by avoiding certain risks to our bodies by staying away from trolley tracks, or avoiding certain donation practices which might make us unable to lead the lives we have planned – seems like something that could be morally significant.

The types of justification offered here do not seem to explain all cases. The *Switches and Skates* judgements, supposedly changing depending upon the addition or removal of alternatives, remain difficult to explain. However, the finding that sometimes our intuitions are mistaken is not a revelation. The claim Unger needs to establish is that our intuitions are systematically mistaken in ways that have no moral relevance. If what I have suggested here is correct, projective separation (and even some protophysical thinking) might be morally justified.

### **Smuggling in Utilitarianism**

Like Peter Singer's *Sacrifice* principle, Unger's Liberationism also seems to sneak in utilitarian thinking. For example, it assumes that our obligations should track values.<sup>32</sup> To elaborate, Unger's contention that our obligations are much more extensive in *Envelope* is based on two claims, the first explicit, and the second unstated:

1. The badness that results in *Envelope*, where we typically do not intuit that an agent has an obligation to act, is much worse than in many cases where we do think acting is obligatory. The differences in these cases do not reflect our true values (Basic Moral Values).
2. If, according to our values, one situation is axiologically worse than another, our obligations should reflect this.

This second claim, however, is not typically borne out in everyday moral thought. If we consider something like the doctrine of acts and omissions – which is accepted in 'common-sense morality' – this does not seem able to cater for our values having such a direct impact upon our obligations.

---

<sup>32</sup>In chapter 6 I will suggest further that this might be a problem.

Something like the second of these claims is required by the Weak Principle of Ethical Integrity, a principle Unger sees as obvious and necessitated by living a morally decent life:

**The Weak Principle of Ethical Integrity** Other things being even nearly equal, if it's all right for you to impose losses on others with the result that there's a significant lessening in the serious losses suffered by others overall, then, if you're to avoid doing what's seriously wrong, you *can't fail to impose much lesser losses on yourself*, nor can you fail to accept such lesser losses, when the result's a *much more* significant lessening of such serious losses overall.

(Unger, 1996, p.140)<sup>33</sup>

Frances Kamm simplifies this as the claim that we “should do to ourselves what we should do to others (or less) to stop greater harm” (1999, p.304). It implies that, given a situation where I can impose some losses on some people to prevent great suffering, we can ascertain that I ought to do less or equal harm (less or equal negative value) to myself if by doing so I can prevent greater harm (resulting in more value). Obviously, an act utilitarian will instantly sign up to this principle. However, even this does not seem required typically under ordinary morality. Kamm notes that some “self-other asymmetry might be sustainable, preventing us from assimilating what we should do to others with what we should do to ourselves” (1999, p.305). There are clearly many asymmetries we readily accept between what an agent may do to themselves and to others. For example, in any ordinary case, it seems permissible for me to arbitrarily impose harms on myself, destroy my property or make bad choices for myself. Realising that I may enjoy a cup of tea more, I might opt for a cup of coffee nonetheless. This would be strange, but certainly permissible. However, if I was to arbitrarily harm another, knowingly make bad choices for them or destroy their property, this typically would be wrong.

Kwame Anthony Appiah makes a similar observation about our commitments when we indulge in enjoyable pursuits rather than contributing to

---

<sup>33</sup>We might suppose that the caveat “Other things being even nearly equal” does a lot of work here, and thus that, as a variety of non-utility features could factor into this, it does not smuggle in utilitarianism as I have suggested. However, even Unger, upon clarifying this, talks in terms of “losses” and “costs”, seemingly presuming that these are the real moral currency in play (1996, p.140).

save lives of distant strangers. We need not say (as Unger suggests we must) that ‘the lives of the children you could have saved were just worth less than your evening at the ballet’ (2007, p.166). Appiah’s contention is that the utilitarian axiology is mistaken; that there is no “single thin currency of goodness” but instead a pluralism of values that cannot be measured the same way. Even if this value pluralism is rejected, it is another question how value relates to obligations. That is, if we are not committed to a maximising utilitarianism, we might fully admit that some course of action is *best*, but suggest that it is still an open question whether or not it is obligatory.<sup>34</sup>

### Rejecting Unger

That our intuitions are susceptible to various strange influences is an interesting finding. However, Unger’s argument that his Liberationism is the only solution fails. I have suggested that the principles Unger notes might not be quite as arbitrary as it first appears. Clearly, however, these features will seem morally distorting to those of utilitarian leanings. Despite attempting to rule out the moral relevance of a considerable number of factors – physical proximity, social proximity, unique or multiple potential saviours, the role of governments, urgency, and causal amorphousness among others – Unger does not consider the potential moral significance of our valuing certain features about our lives.

Ultimately, the way Unger dismisses these ‘distorting’ biases also leads him towards an extremist view. These ‘biases’ often have the effect of justifying some doing/allowing distinction. In dismissing these as morally irrelevant (which, as I have already suggested, is perhaps dubious), this prompts a more utilitarian evaluation where costs are simply weighed. If we reject this manoeuvre, the extremist conclusion no longer follows.

## 2.4 Dealing with Extremists

Singer and Unger are paradigmatic examples of philosophers who would dismiss demandingness objections, because they would accept that morality is extremely demanding. I have argued that their efforts to demonstrate this do not succeed. However, there are several other ways one might attempt

---

<sup>34</sup>This type of thought is developed in chapter 6.4.

to disregard demandingness objections. In the remainder of this chapter, I briefly discuss and offer responses to some of these.<sup>35</sup>

### 2.4.1 Sidgwick

I intimated at the beginning of this chapter that common-sense seems to be on the side of undemanding theories. However, Sidgwick is an example of an extremist, but claims that common-sense is on his side:

“Certainly we should agree that a truly moral man cannot say to himself, “This is the best thing on the whole for me to do, but yet it is not my duty to do it though it is in my power”: this would certainly seem to common sense an immoral paradox.”

(Sidgwick, 1907, p.220)

I say Sidgwick is an extremist (or at least that his position seems to entail extremism) because it seems obvious that requiring *the best* all the time will lead to extreme demands (at least right now). This claim of his does not seem to resemble common-sense. It was noted earlier how odd it would seem to say that someone donating large portions of their wealth (but not *as much as* she can) and regularly volunteering (but not *as much as* she can) is morally failing. We see people like this as extremely admirable. One reason this seems so bizarre is that an evaluation of moral failure seems to imply *blame*.<sup>36</sup>

Singer and de Lazari-Rabek interpret Sidgwick as defending a version of utilitarianism wherein we separate the connection between blame and wrongness. This might then alleviate the concern of *blaming* someone who does not act as well as she can. She would be doing something *wrong*, but might not be blameworthy. Blame, under this picture, is only appropriate when the process of blaming would have good consequences. Singer and de Lazari-Rabek give an example of a mother who prioritises her child over others because of her affection for him (2014, p.332). The love she has for

---

<sup>35</sup>One type of extremist I do not engage with here is the sort, like Lisa Tessman, who deny the ‘ought implies can’ principle, and accept that there are cases of impossible moral requirements (2015). Several of the criticisms in this section will be applicable to Tessman’s view, but with regards to Tessman’s position in particular, I agree with Carbonell’s assessment that we lose a lot if we do not view moral requirement as “action-guiding, something we can really hold you to” (2015).

<sup>36</sup>Blame is discussed much more extensively in chapter 5.

her child which causes this behaviour, they suggest, should be promoted (and not condemned) even if it sometimes leads to acts which are morally wrong.

Sidgwick does explicitly endorse the praising of “any conduct more felicitous in its tendency than what an average man would do under the given circumstances” (Sidgwick, 1907, p.493). This would allow for us to praise the regular donor and volunteer mentioned earlier, and maintain many of our similar praising/blaming practices.

However, I am unsure whether this manoeuvre is successful. I have two concerns:

- 1 They seem to advocate an extremely functionalist ideal of blame – we should only blame behaviour so long as we want to discourage it and should only praise behaviour to increase its frequency. While this no doubt is an aspect of what it is to blame, this does not seem to capture the full picture.
- 2 Once wrongness is separated from blame in this way, what remains? Wrong acts are typically taken to be those which are forbidden. Not only would this entail that many wrong actions should not result in blame, but it does not even seem appropriate that the agent committing the so-called ‘wrong’ – such as a mother prioritising her child – should feel any moral regret. After all, they may have acted on motivations which typically produce the most good, so they should continue to act in this way.

If blame takes this back-seat, its use being appropriate when (and only when) it will cause the best consequences, we leave conceptual space open for its being appropriate to blame even those performing heroic acts. If it so happens that doing so will have optimistic consequences, blaming (and even punishing) someone acting perfectly well will actually be obligatory.<sup>37</sup> Hooker also notes that this clashes with our intuitions, which “hold that whether or not someone is to blame or is to be punished is not really a matter of whether on this occasion blame or punishment would produce good

---

<sup>37</sup>This point is made by Norcross (2006, p.226), who rejects the notions of rightness and wrongness altogether. Singer and de Lazari-Radek also discuss cases where one ought to be *punished* – or at least “encounter a hostile public reaction” – for doing *the right thing*, when doing so involves breaching rules that are typically useful (2014, p.309).

consequences but instead a matter of whether the blame or punishment is warranted by the person's conduct" (2016, p.145). Considering Sidgwick's desire to preserve a common-sense view of morality, this seems like a major flaw. It seems absolutely bizarre to suggest that practically everybody is so mistaken about when blame and punishment are appropriate.

If we hold that coherence with our intuitions is to be of any guide in determining what an appropriate theory should look like, a theory with repercussions this incongruent with our practices must be seen as problematic. Given that Sidgwick saw common-sense as a useful guide, this approach fails on his terms.

#### 2.4.2 Denying Overridingness

In the previous chapter, it was briefly mentioned that another way to respond to an overly demanding theory is to accept that morality really is extremely demanding, but add that there are other types of consideration that matter when evaluating behaviour generally. People like Dale Dorsey endorse this position, which has the upshot of enabling us to accept the arguments given by the extremists, but deny — as seems desirable — that there is anything irrational about not meeting the apparent moral obligations these accounts.

The general thought is that morality is only one part of human life, and not the ultimate arbiter of behaviour. Something being morally best is only *one consideration* in its favour, and needs to be weighed against other considerations. When considering a theory like maximising act utilitarianism, for instance, someone denying overridingness can claim that acts that are not *the best* in terms of utility are morally wrong, but they are not *all things considered* wrong.

Demandingness objections claim that if a theory is too morally demanding, it should be rejected. The move here suggests that moral demands are merely one constituent consideration of how we should evaluate conduct. We might say that morally you ought to  $\phi$ , but *all-things considered*, you need not.

Dale Dorsey describes a way that this can account for supererogation (2013). In his picture, just like acts might be *morally* forbidden, permissible or required, they may also be *rationally* forbidden, permissible or required. Because he denies that moral reasons have 'supremacy' over other reasons,

like prudential reasons, these compete. Moral considerations are only one thing to consider among many factors which affect one's decision. If the moral reasons are not strong enough, they lose out, and the morally best act is not required. Supererogatory acts are then understood as not "better than is morally required, but better than is rationally required" (2013, p.373).

While this provides a neat picture, it does not allow us to properly appreciate a moral failing as something that particularly warrants the judgements we expect in those cases. When someone does something that is prudentially good – i.e. good for them – but bad for others, this merits a special type of criticism.

Dorsey's account also seems to run the risk of distorting our moral categories. Consider for instance, a 'Good Samaritan' type scenario – you walk down a country road and see a traveller in need of urgent medical attention. Common sense might give us verdicts like the following. Just walking by and ignoring them (or robbing them!) is forbidden. Helping them to a nearby hospital, or calling for help, might be required. But it seems that in addition to securing medical attention, offering them money or some of your clothes to make them more comfortable (embellish the case as you need to to arrive at the judgement that they do *more* than they had to) is supererogatory. If acts were morally graded, supererogatory acts are those that get a figurative 'gold star'.

According to Dorsey's picture (particularly when combined, as he thinks desirable, with a maximising consequentialism) even those middle options (which were just good enough), if they are better than rationally required, must also be evaluated as supererogatory.<sup>38</sup> This demeans the category, which is understood as including truly *optional* acts; not those which seem clearly required. In short, a picture that denies overridingness fails to capture our rich moral landscape.

---

<sup>38</sup>Dorsey notes an objection of this sort, that his account will "implausibly extend the category of the supererogatory", but does not think it a problem (2013, p.377). He does supplement the view with some additional fixtures that make it more conducive to our intuitions – e.g. making it a condition that an act is significantly morally better than other options in order to count as morally good, which is a condition for being supererogatory – but these do still leave strange categorisations, which he acknowledges.

### 2.4.3 Evaluation

In this chapter, I have discussed several prominent arguments in favour of extremism. If extremism is correct, the demandingness objection is a bad objection – morality should make extreme demands. I have argued that we should not accept extremism, favouring instead the moderate position. In the remaining chapters, I do not argue any more for this position, though I make a concession to the extremists in the final chapter, suggesting that this type of investigation informs us about what they should say about moderates and the intuitions which support a moderate position.

Accepting that the extremists are mistaken entails that the demandingness objection does provide us motivation to reject overly demanding moral theories. We may wonder what it is about those theories that causes them to be excessively demanding. I consider several options for diagnosing the problem in the next chapter.

## 3

# What Makes a Theory Overly Demanding?

*“Where did it all go wrong?”*

– Oasis, *Standing on the Shoulders of Giants*

There are several obvious suggestions we might consider for what makes a theory overly demanding. Here I shall consider these, as well as some less obvious candidate properties and discuss whether they have any structural role in generating demandingness objections.

### 3.0.1 Aside: Types of Cases

Before discussing the features of a theory which make it vulnerable to demandingness objections, it might first be worth considering the types of examples this objection yields and what makes the verdict seem too demanding in each case. The cases given will all be question-begging to an extent. For example, the type of case given most so far – an iterative case like Cullity’s – when taken as a paradigm example of a problem case will claim that the situation is too demanding. This will of course be rejected by those who see the objection as misplaced because they think morality actually is (and should be) extremely demanding.

The main type of case that will be of use is the iterative case. Any theories which allow for plausibly small demands to be iterated repeatedly until an extremely high level of demand is reached will face charges of demandingness.

On top of this, we might consider cases that give one-off verdicts which seem too demanding. We must be careful here, however, to ensure that what is objected to in the case is the *demandingness* of the theory, rather than some other perceived issue it presents. To illustrate this, we might consider van Ackeren and Sticker’s example of demandingness in Kant:

“The specific form of Kantian demandingness [they discuss] comes in the form of the threat that an agent at any time might find herself in a situation in which she has to sacrifice all of her non-moral goods. Such a case is explicitly described in the *Second Critique*’s gallows case in which an agent is rationally required to sacrifice his life (he is sent to the gallows) when a tyrant demands of him on pain of execution that he lies in court.”

(2015, p.85)

In cases like this, where it is suggested that one is not permitted to lie even about something trivial – and even when doing so might save one’s life (or even the lives of many others) – it might strike us that the reason this is so wrong is not its *demandingness*, but that it gives the wrong *verdict* of what the agent should do.<sup>1</sup> Of course, if one did think telling the truth was the right (or best) action in this case, but thinks it too demanding to be obligatory, this would be the type of case that would provide a demandingness objection.

To give a hopefully clearer example, we might think that there are single actions that might be accepted to be the best action available, but such that we do not think it compulsory because of the demand it would impose upon the agent. Suppose that a nearby rescue situation arises and the agent with the ability to save the victim estimates that there is a 51% chance she could save him, but a 49% chance she would fail and die in doing so. We might think we should not blame her for not putting herself at such a risk, even if it appears to be the best thing to do, and that this indicates that failing to perform the best act is not morally wrong. A theory that made such acts requirements would probably appear overly demanding.

---

<sup>1</sup>McElwee discusses objecting to a theory on these grounds. He labels this the “The wrong moral ranking objection” (2017, p.89), which he notes is logically distinct from the *pure* demandingness objection.

Similarly, we might consider how demanding it could be to have some acts forbidden. In Ashford’s discussion of Scanlon’s contractualism, for example, she suggests that in any practically realisable state of the world, personal air travel will be forbidden (2003, p.299). This is because in Scanlon’s account principles permitting actions which have a chance — even a very remote one — of being harmful may be reasonably rejected by agents who cannot benefit (or expect to benefit) from those actions (1998, p.208). Because personal travel can only be expected to be enjoyed by the very rich, the slight chance that those planes could crash into persons or property who could not benefit from them is enough for a reasonable rejection. We might think that a theory which contains requirements to abstain from certain individual acts could also be overly demanding.

Having given types of cases that are iterative act-requirements, singular act-requirements and singular restriction-requirements, a fourth category seems to present itself: iterative restriction-requirements. For such an example, we might simply think of cases like the above — say a prohibition against certain modes of travel — but where it seems an individual instance would not be extremely demanding. Not making use of a car/plane (or something else useful) on one situation, it might be supposed, would not be overly demanding, but if this is repeated ad infinitum, it may impose significant demands. This is sometimes described as a confinement objection.<sup>2</sup>

It seems that a theory might be seen as overly demanding if it can generate any of these types of cases. With this in mind, I will now consider various features of theories that might make a theory susceptible to this.

### 3.1 Obligation-Generation

We may begin with a fairly obvious condition. As, under the analysis of the demandingness objection I provided, the objection claims of a theory that it “gives a verdict which demands significantly more” than we would otherwise expect, the theory must demand *something*. It must actually make demands. Thus an obligation-free moral theory is immune to demandingness objections.

We might wonder whether removing obligation-generation from a moral theory is a tenable manoeuvre, as this seems like such a significant part of

---

<sup>2</sup>e.g. Benn (2016).

what a moral theory does. If a moral theory cannot tell us what we ought to do in a situation, it raises the question of what we want a moral theory to actually do. Elizabeth Anscombe, however, in “Modern Moral Philosophy” claims the notions of “duty” and “obligation” are damaging to moral philosophy (1958, p.1). Thoughts of this nature contributed to the development of contemporary virtue ethics, but there are other theories which incorporate this move, even among consequentialists. Alastair Norcross’s scalar consequentialism, for example, also avoids the notion of obligation. Norcross properly observes that rejecting “the notion that morality issues demands at all” completely avoids the demandingness objection, and suggests that this is an attractive option (Norcross, 2006, p.219).

Moral theories that lack an account of obligation, or deny that obligation is a useful notion, must account for our moral practices and phenomenology, as well as explaining quite how the theory can be action-guiding. Our moral practices of blaming seem appropriate if an agent was morally obligated to perform an action and fails to do so. An agent is at fault because they failed to meet the requirements of them; for not doing what they *had to* do. The moral phenomenology – feeling that one is free to perform or not perform an action in some circumstances, but that one is compelled to do so in others – also seems to cohere sensibly with the notion of moral obligation. We might well think, as Robert Goodin claimed, that a “morality that was never prepared to make any demands would be a pretty useless morality” (2009, p.3).

While I acknowledge that entirely removing the concept of obligation from a moral theory is a possible manoeuvre, which completely avoids the demandingness objection, and is seen as desirable by some, I will not discuss such theories any further here. All that must be said here is that though obligation-free moral theories avoid this criticism, they have difficulties (of the sort given above) of their own. From here, I will focus entirely on theories which do accept and invoke some notion of obligation.

Even accepting that a theory’s being obligation-generating is required for the availability of a demandingness objection, we may still ask just what type of obligations might be generated, and whether a theory must generate a specific type of obligations. Imperfect duties in Kant’s moral philosophy, for example, may be discharged in a variety of ways. This latitude or ‘play-room’ with regards to deciding how one might act – and still be acting

morally – may affect how demanding a theory might be.<sup>3</sup> When considering one’s duties of beneficence, it might be less demanding if the demands are unconfined, for example, if one may choose between spending time with the elderly, helping at a homeless shelter or donating to a charity.

Kant’s view involves both imperfect and perfect duties. Perfect duties, like the duty not to steal, permit no exceptions. They “prescribe the actions we are to take (or omit)” (Baron, 1995, p.30). Imperfect duties like beneficence may be discharged in perhaps any of a large disjunction of ways so can be seen as coarse-grained duties. This might be contrasted with act-utilitarianism, which we might see as typically generating only fine-grained obligations.<sup>4</sup> It seems however that whether obligations are fine-grained or coarse-grained, they could lead to overly demanding prescriptions. It is important for Kantian imperfect duties that they cannot simply be ignored; they must be discharged. So, we may imagine that even with a long disjunction of possible ways to discharge a duty, each way would involve something overly demanding. Alternatively, we might look at the iteration of imperfect duties. If discharging the duty required performing several actions, each individual one might not be demanding, but the series as a whole might be.

So obligation-generation (whether the obligations be coarse-grained or fine-grained) is a necessary condition. It is clearly not sufficient, however, as a theory could generate only minimal and occasional demands.

### Maximisation

In the previous section, it was established that in order to be too demanding, a theory must make demands. This does seem trivially true. We may wonder whether it guarantees the availability of a demandingness objection if these obligations are taken to be maximising.<sup>5</sup> McElwee suggests that one

---

<sup>3</sup>This is likely to depend on the interpretation of Kant given. Some authors, like Thomas Hill, suggest that imperfect duties are where we may locate the category of supererogation in Kantian ethics. Hill claims that duties of beneficence are the duties of ‘widest’ obligation thus allow most latitude (2002, p.204). Other Kantians, such as Marcia Baron (e.g. 1987; 1995) and Jens Timmermann (e.g. 2005) deny that Kant makes room for supererogation. As a result, any latitude under their interpretation might not affect the demandingness of an agent’s obligations.

<sup>4</sup>I say *typically* because there may be situations of a tie in terms of expected utility, and such a tie could be among a large number of options. In those situations, a disjunction of options equal to the number in the tie would be available.

<sup>5</sup>The term ‘maximising’ makes more sense when considering utilitarian theories, where some good (utility) is maximised. I will continue to use this term, but with regards to a

reason consequentialism is seen as so vulnerable to accusations of over-demandingness is that in “its modern versions, consequentialism has generally included what we might call a presumption of optimality – the idea that ‘only the best is good enough’” (2017, p.86).

For any values a theory may esteem, it might be suggested that an agent should – and is morally required to – promote those values to the highest degree she can, or act in accordance with them at all times. It is often supposed that utilitarianism ordinarily requires that an agent maximises the amount of happiness, and, as we have seen, this results in an extremely demanding account of morality. If maximising would have a similar effect for all plausible conceptions of the good, this would suggest that maximisation is the problematic feature.

Some types of obligations, however, are not demanding even when an agent acts in accordance with them as much as possible. An obvious example of this would be some negative obligations – those following from negative duties. If we have the (plausible) obligation not to gratuitously torture<sup>6</sup> anyone, maximally cohering to this obligation – never gratuitously torturing anyone – is actually a very easy requirement. If the values of a moral theory only required actions like this, presumably they could all be compatible with a very easy lifestyle.

From this we might want to conclude that it is not sufficient for a theory to be maximising in order for a demandingness objection to be available, that maximising negative duties need not be that demanding. However, this would be too quick. Any theory that is to be plausible must place some value on the wellbeing of others (though perhaps not unconditionally, e.g. Kant). This should be a value that, in the right contexts, could confer obligations upon a moral agent. In a case like Singer’s pond example, or to take an even more extreme case, if one could save a million people simply by pushing a button, it seems like a theory would be amiss if it did not give the verdict that the agent was morally obligated to act.

This may have been a little brief, but for reasons like these it does not seem controversial to claim that a moral theory should (in the right circum-

---

variety of other theories, an obligation to *do the best* is more accurate.

<sup>6</sup>I include the term “gratuitously” to cover cases where torture might be argued to be important (e.g. with the goal of saving many, many lives), and strange situations wherein not being able to could be claimed to be demanding. I do not want to take a stance on that question here, but it seems that *gratuitous* torture must be easy to avoid.

stances) entail positive obligations. Any theory that does value the wellbeing of others and requires maximisation of wellbeing, will be extremely demanding. An alternative option is that a theory may claim that maximising the good is obligatory, but have the wellbeing of others as only a component in what ‘the good’ amounts to. On a Kantian picture, for example, as well as duties to promote the welfare of others (duties of beneficence), an agent has duties to oneself. Some of the duties to oneself are *perfect* duties too, like the duty not to commit suicide, or “deprive oneself of an integral part or organ” (*Doctrine of Virtue*, [§6:423]),<sup>7</sup> or eating or drinking to the degree of incapacitation (*Doctrine of Virtue*, [§6:427]). We might think that these perfect duties to oneself could limit the amount of beneficent conduct that could be required of an agent.

Maximally promoting the good (as much as this makes sense) in Kantian terms, prohibits violation of any perfect duties (as that would make an action wrong) and might require taking some action to take other imperfect duties seriously. Some of these imperfect duties are also to oneself, such as the duty of an agent to “cultivate his capacities” (*Doctrine of Virtue*, [§6:445]). If we judge Kant to require that an agent make some time or resources available for cultivating one’s capacities, even at the expense of potentially life-saving beneficent acts, we might wonder if this would reduce those commitments enough to not be susceptible to a challenge of being overly demanding. While a Singer-esque utilitarian will suggest that one must help others as much as possible until doing so will involve more sacrifice for oneself than other agents would benefit, a Kantian might claim that acting like this would force them to neglect their talents. Pinheiro Walla suggests that one motivation for Kantians to incorporate imperfect duties is to prevent morality being “excessively demanding, perhaps even unbearable for human beings” (2015, p.732). Prima facie, accepting that there are occasions wherein one has a choice whether or not to perform some beneficent act, and accepting that one has moral obligations to oneself that compete with obligations to others, does seem promising for avoiding over-demanding prescriptions.

Because of the apparent promise of manoeuvres like this, we may ask whether an acceptance of multiple sources of obligations – still accepting

---

<sup>7</sup>Kant distinguishes here between integral parts or organs on one hand, and diseased organs or mere “parts” on the other. Through this distinction he allows that cutting one’s hair is permissible, but giving away a tooth is maiming oneself, and is regarded as “partially murdering oneself”.

that an agent must always do the best – exempts a moral theory from overdemandingness. My response to this is two-fold.

Firstly, in Kant’s case specifically, it seems like even the perfect duties to oneself might incur extreme (and implausibly high) demands. Consider the perfect duty not to deprive “oneself of certain integral, organic parts” (*Doctrine of Virtue*: [§6:421]). This could become extremely demanding if an agent could save the life of a loved one only by donating a kidney or part of one’s liver. Being forced to watch a loved one die perhaps from a condition that entails a great deal of suffering – due to obeying this perfect duty – would be extremely demanding.<sup>8</sup> Pinheiro Walla offers several other cases, where the supposed normative priority of perfect duties over imperfect ones leads to counterintuitive verdicts, such as the obligation not to violate a person’s private property (a perfect duty) when doing so would be necessary to save someone’s life (falling under the duty of beneficence – an imperfect duty). In addition, she notes that even just complying with perfect duties, such as the duty to keep promises, can when there are “unfriendly external circumstances”, be extremely demanding (2015, p.741). Tyrants demanding that one lies or face threats of severe punishment could be one such example.

Second, we can think of another way a theory may include maximising and include duties that could compete with the duty to promote welfare. Consider W. D. Ross, for instance, who argues that there are several sources of duties. As well as duties of beneficence, Ross identifies duties of fidelity, duties of reparation, duties of gratitude, duties of justice, duties of self-improvement and duties of non-maleficence (2002, p.21). It would be possible to have a theoretic weighting to these duties, suggesting that acting upon the ‘heaviest’ duty in a given situation is obligatory. An agent in this situation is still required to perform *the best* course of action, but this will not necessarily result in the greatest happiness or definitively prohibit actions like promise-breaking.

On a Rossian picture, it is not obvious what maximising might look like, but it seems plausible that sometimes the ‘best’ thing to do, whatever that might involve, will be hugely demanding for moral agents.<sup>9</sup> In iterative

---

<sup>8</sup>As in the case provided by van Ackeren and Sticker, we might think that this case is another instance where the problem is not that the act is too demanding, but that Kant arrives at the wrong verdict here.

<sup>9</sup>This point is also noted by Liam Murphy (2000, p.10). He suggests that Ross – as well as Samuel Clarke and perhaps William Godwin – advocates a pluralism which

examples, for instance, it still seems that a great deal will be required (even if not maximising the amount of lives saved).

However a theory accounts for valuing wellbeing, it seems that maximising is at least likely to lead to severe demands. If wellbeing is seen as the only value, maximising leads to maximising utilitarianism, which is extremely demanding. If wellbeing is valued among other values, any theory which maximises all of these will at least have a difficult time explaining how maximising these values is not extremely (and excessively) demanding.

It clearly seems that obligations do not have to always require doing the best in order to be overly demanding. To illustrate this, we might consider a form of satisficing consequentialism. Satisficing consequentialism claims that agents “may sometimes choose what is good enough, without regard for whether what they have chosen is the best thing (outcome) available in the circumstances” (Slote, 1984, p.141). If we consider a ranking of actions from morally best to morally worst, a satisficing theory claims that of a certain act, an act at least that good must be done, while an action lower on the ranking would be impermissible. Higher acts on the ranking may be supererogatory. If a satisficing account places the bar very high – but still not maximally high – it might still seem too demanding, and lend itself to a demandingness objection.

Thus, a (plausible) theory’s being maximising is definitely not necessary for a demandingness objection to be applicable. Whether, combined with a plausible account of values, a theory’s requiring maximising is actually sufficient for the objection, I am not sure, but would tentatively suggest that there is good reason to think it is.

Having established that obligation-generation is a necessary feature for a theory to be susceptible to demandingness objections, no other features could be sufficient, absent obligation-generation. Other features might still be necessary, or they may be jointly sufficient alongside obligation-generation (and possibly others). Consequently, for the remainder of features I consider whether they are necessary or jointly-sufficient with obligation-generation.

---

involves a maximising (or “optimizing” in Murphy’s terminology) principle of beneficence (2000, p.137,n.3), and that this is likely to lead to very demanding accounts of morality.

## 3.2 Impartiality

One feature of moral theories that we might think is likely to make them very demanding is moral impartiality. If we accept that everyone is of equal moral value, then we might be persuaded that the suffering of oneself or a loved one should have the same ability to generate demands as someone on the other side of the world. If everyone counts equally, and the lives of world's poor are improved much more easily, we might be drawn, *prima facie*, to the conclusion that we (the relatively affluent) ought to devote almost all of our energies to relieving the suffering of the distant poor. It might seem impartiality is bound to lead to excessive demandingness. Ashford, for example, suggests that “in the current state of the world it may not be possible to defend less demanding obligations to those in need within an impartial framework” (2003, p.274).<sup>10</sup>

That we might ordinarily think impartiality to be important in generating the objection is noted by Garrett Cullity, who claims that problems of demandingness have two conventional theoretical sources noted in the literature, and that they both have in common the claim of impartiality. He suggests that once an agent accepts she is no more important, morally speaking, than any other agent “it is easy to see how an extremely demanding morality might result” (Cullity, 2004, p.92).

One way the problem is ordinarily generated is through impartiality in relation to what is deemed to be valuable. If the moral evaluation of an action is determined by the value of an outcome, and the evaluation is agent-neutral, then a simple method of generating obligations (such as the maximising criterion of maximising act utilitarianism) will result in very burdensome demands.

A second way Cullity notes that extreme demands may result from impartiality is through responses to reasons (Cullity, 2004, p.92). If an agent accepts that her suffering (or desires/hopes/dreams) gives her a reason to act, and that due to impartiality, the force of the reason to act should be the same regardless of whose suffering grounds it, she is committed to the demanding claim that she should respond to the needs of others with the same vigour as she would her own.

Despite this, Cullity rejects that impartiality should be seen as essen-

---

<sup>10</sup>A similar point is made in (Ashford and Mulgan, 2012, §8).

tial to generating the problem. He notes that iterative methods, like those described in 1.2.1 can generate excessive demands, and relies on neither of these sources. This can be demonstrated by generating with a very partial theory. Cullity arrives at extreme demands by considering a low-cost life-saving scenario and considering multiple iterations of it. Even if the theory does involve some partiality, if it still requires an agent to sacrifice something negligible in order to save a life in some case, we can iterate the obligation until it does become extreme. Problems of demandingness, Cullity argues, arise even if we were to accept the following, Partial Premise:

“It may be wrong for me to do nothing to help another person, but it is not wrong to be prepared to do significantly less for him than I would do for myself.”

(2004, p.93)

This premise clearly relies on neither of the traditional sources of excessive demandingness that Cullity mentioned. Cullity uses this to cast doubt on responses to extreme demandingness which try to justify a claim like the Partial Premise (like Scheffler’s agent-centred prerogatives do), and escape that way, as this is not enough to actually escape the problem. As the demandingness objection can be generated even if a theory denies impartiality, it cannot be a necessary condition for a theory to fall prey to the objection.

When we consider impartiality paired with obligation-generation, it is also not jointly sufficient. That is to say, obligation-generating theories which are also impartial may not be susceptible to the objection. There are no obvious examples (to my mind at least) of impartial, obligation-generating theories which do not fall victim to the objection.<sup>11</sup> Even a theory which makes extremely small demands can result – through several iterations – in extreme (and implausible) demands, so it is difficult for a theory to avoid the possibility of extreme demands.

However, we can manufacture theories to meet this condition. Consider Defenestration-Empathy-Theory (DET). This bizarre moral theory acknowledges that all people are morally equal, but holds that moral requirements only obtain when someone is thrown out of a window, and they only require

---

<sup>11</sup>Miller (2004), Cullity (2004) and Murphy (2000) attempt to provide accounts that meet this condition in response to concerns of demandingness. Miller’s and Murphy’s accounts are discussed in chapter 4.

a display of some empathy with the defenestrated agent. Clearly, this is not a theory many will rush to endorse, but it is difficult to see how a demandingness objection could apply to such an account. Even if iterated to a large degree – if you happen to be surrounded by people who are falling out of windows – displaying some empathy to each of them does not seem like it could be too demanding. Our actual obligations – ones we are generally willing to acknowledge – will be heavier than those issued by DET.

We might respond to this finding by suggesting that this is not what we want *impartiality* to mean when talking about a moral theory. Bernard Gert defines impartiality the following way:

“A is impartial in respect R with regard to group G if and only if A’s actions in respect R are not influenced by which member(s) of G benefit or are harmed by these actions.”

(Gert, 2004, p.117)

This might be thought to imply that impartiality required equal treatment. If we consider the generosity, for example, to be impartially generous, one would need to be generous in ways that are not influenced by who one is being generous to. Buying a gift for a friend – because they are a friend – would then seem problematic, as this would clearly be influenced by who would benefit from the act. We generally do not think that theories claiming to be impartial make this kind of commitment, so this unrestricted impartiality is not required for a theory to be considered impartial – a theory’s being impartial does not require that the beneficiary of every action which harms or benefits to be impartially determined.

There are several candidates that might be intended. Hooker discusses three things that might be meant by impartiality in “When Is Impartiality Appropriate?”. These are:

1. Impartial application of good (first-order) moral rules.
2. Impartial benevolence as *the* direct guide to decisions about what to do.
3. Impartial assessment of (first-order) moral rules.

(Hooker, 2009b, p.26)

Requiring impartial application of good moral rules is more subtle than simply requiring equal treatment of everyone. A good moral rule may (very plausibly) allow a person to treat certain groups differently. Many rules – particularly those concerning violations – can often easily be impartially complied with, and doing so involves no partial behaviour. The requirement not to kill, for example, must be impartially applied, and this is done by simply not killing anyone (Gert, 2004, p.118). Some rules however will mandate partiality. Sidgwick, for example, defends the partial rules that utilitarianism might result in, as it would “clearly not promote the universal happiness for each one practically to concern himself with the happiness of others as much as with his own” (1907, p.431). The rules that result then may allow, or require, a person to treat themselves or other groups differently. We might consider the rule:

“The product of a collective enterprise should be divided in proportion to people’s contributions to the enterprise”

(Hooker, 2009b, p.29)

If this is a good rule (and it looks very plausible), it will often result in dividing the products of the enterprise unevenly. How to apply a rule might still be complicated. If a band agrees to implement a rule like this, for example, they might consider who purchased music software, booked gigs or provided practice space when applying the rule. It might be that a fair application of a rule like this would result in one person receiving considerably more of the product of the enterprise, and this would not be a problem. When the rule is applied, however, it must be done so impartially. If Ringo booked the gigs, for example, this should not be seen as a less valuable investment in the enterprise than if Paul booked them.

A similar example might be the rule that “Employers should pay their staff what they are entitled to”. We might imagine that most employers do agree with this rule, yet it still might be applied partially. In workplaces where women or ethnic minority workers are paid less than their white male counterparts, for example, it might be suspected that some partial explanations are used in the application.

Hooker correctly claims that impartiality in the application of good moral rules is “always appropriate” (2009b, p.26). It seems that this should

be required of any impartial moral theory. However, this understanding of impartiality does not seem to generate overly demanding moral verdicts unless the first-order moral rules themselves are overly demanding.<sup>12</sup> A theory could also be extremely demanding even if its rules are not applied impartially. We might imagine a set of good rules – ones we might be inclined to agree with – but under a meta-ethical principle that we only apply the rules partially to members of our own race. Because members of one’s own race may require large amounts of aid, the racism of this meta-ethical principle would also not prohibit extreme demands. While we do accept that impartial application of rules is important, this is not typically what we mean thinking of a requirement for impartiality.

The second impartiality claim Hooker discusses – impartial benevolence as the direct guide to decisions about what to do – concerns completely unrestricted impartiality about reasons to act. It suggests that benefits “to strangers would count in your reasoning no less than the same-sized benefits to your partner, child or mother” (2009b, p.31). Theories that standardly require for type of impartiality would certainly be extremely demanding. Hooker notes that if an agent is able to save three lives by sacrificing herself and donating her heart and both kidneys, this would be required (2009b, p.31). Maximising act-utilitarianism may make these sorts of requirements, but it is not seen as a requirement for a theory’s being impartial. One type of case where impartial benevolence is required as a guide is in certain official capacities. People running charities, for example, who are charged with catering to the welfare of famine-ridden countries are obliged in their professional capacity to use the resources at their disposal completely impartially.

If, for a theory to be impartial, it needs to require impartial benevolence as the guide as to what to do (which I have suggested is not the case), then this does result in what seem like excessive demands. We know this from the example of act utilitarianism. This is because as soon as an account of welfare is plugged in – to say what goods or harms consist in – it would require an agent to treat the goods or harms of any other persons as just

---

<sup>12</sup>A strange feature of this particular analysis is that it allows most theories to be cast as impartial, so long as the rules are accepted as good rules. Scheffler’s agent-centred prerogatives, for example, could be defended from an impartial position. Even forms of egoism – perhaps Brink’s rational egoism (1997) – might be regarded as impartial, despite being, what we understand as paradigmatically partial accounts.

as obligation-generating as those of herself or loved ones.<sup>13</sup> We can also note that a theory does not need to be impartial in this sense in order to be overly demanding, as is demonstrated by several accounts already discussed which can be overly demanding but are not impartial in this way, such as Scheffler’s agent-centred prerogatives.<sup>14</sup>

The third of Hooker’s considerations is what I take to be definitive of what a theory needs in order to be properly considered impartial. It requires that the assessment of which rules that a theory should have must be impartial. What this requires is not as simple as might be suspected at first. It seems that an impartial justification must be one that is “defensible from an agent-neutral point of view” (2009b, p.35). This will then include assessments which permit no partial evaluations, like the utilitarian assessments of actions, which see benefits or harms as having the same importance regardless of who experiences them, but also some that do allow partial rule-formations. Consider rules that might benefit the poorest arrived at through reasoning from behind a Rawlsian veil of ignorance. The rules will be partial, but – if it is accepted that they result from considerations of justice – impartially justified.

Commitments to impartiality of this sort seem essential for a theory to be an impartial one. Unfortunately, it is still difficult to see which rules should be adopted. As Hooker accepts, “no appeal to impartiality by itself can resolve the question of which agent-neutral considerations and distinctions do, and which do not, have foundational importance” (2009b, p.39). Because of this, theories may be justified on impartial grounds and place a heavy focus on property rights and just transactions, like Nozick in *Anarchy, State and Utopia* (1974), or place emphasis on maximising the minimum level of welfare in a society, like Rawls in *A Theory of Justice* (1991). Hooker notes that rules which implausibly benefit those with certain talents, like song-writing or sporting ability, can be agent-neutrally defended, yet lead to highly unpalatable elitism (2009b, p.37).

Construed in terms of impartial assessment of first-order moral rules, impartiality (and obligation-generation) would also not be sufficient to result

---

<sup>13</sup>Given an appropriate state of the world, this would then be very demanding. As mentioned in section 1.1.1 it does not matter whether this state of the world actually or currently obtains.

<sup>14</sup>Other examples include Miller’s *Sympathy* principle, discussed in section 4.1, and Murphy’s Collective Principle of Beneficence, discussed in section 4.2.

in extreme demands. We might think the Defenestration Empathy Theory example could still be an example of this (an impartial theory that is not too demanding) but this might be rejected as not impartially defensible (as it simply does not seem like anyone could really endorse such a theory). A satisficing consequentialism with an extremely low bar for requirement presumably could be defended impartially however, and the conditions for satisficing might be stipulated to be so low that it is never too demanding.<sup>15</sup> This interpretation of impartiality could also not be necessary, as we saw by considering the society with horrifically racist rules (or which make specific reference to individuals) which is also extremely demanding.

When we consider impartiality as a virtue of a moral theory, it involves a requirement not to be *biased* in favour of some group or another. We want our theory to embrace equality in some relevant way. For impartiality to be plausible as a feature of a moral theory, however, we want it to be compatible with considerations of justice. For instance, it must be permissible to reward people who have worked hard, in favour of those who have been idle. This is agent-neutrally defensible, but will result in partial rules. It also seems that we must, in the way Sidgwick discusses, be able act partially (to at least *some* degree) with respect to our loved ones – even if this behaviour should be justified agent-neutrally. It therefore seems that, as far as impartiality *is* something desirable in a moral theory, it is that its rules must be agent-neutrally defensible.<sup>16</sup>

This should demonstrate that a theory’s accepting impartiality is neither necessary nor sufficient (even alongside obligation-generation) for a demandingness objection to be available. Despite often being picked out as a key feature in generating overly demanding verdicts, impartiality actually plays no crucial role in a theory’s being diagnosed as overly demanding.

### 3.3 Insensitivity to Social Factors

There are many ways in which social factors may affect an agent’s phenomenology of moral obligation. By *social factors* I mean the values, expectations and practices of an agent’s community. There is an obvious sense

---

<sup>15</sup>Jan Narveson seems to want to defend an account of our obligations like this, claiming that insofar as we can talk about duties to the faraway needy, they are fairly minimal, and that when people do more than this, it is “nice”, but not obligatory (2003).

<sup>16</sup>I take it as obvious that these rules should also be impartially applied.

in which these factors will affect one's obligations. If, for example, people around you regularly participate in certain dangerous sports, this may affect your obligations in that you will more often face situations where you are able to provide aid. Similarly, if they have certain values, certain behaviours may offend or upset to them to the point that their welfare will be significantly diminished, and we might think that this grounds obligations around such behaviours. Of course, these sorts of examples do seem relevant, but they are not what I am concerned with here. What I am concerned with is ways social factors may affect obligations that are not obviously accounted for by changes in welfare.

In this section I consider three ways in which social factors may affect obligations. First, they may make certain acts easier or more difficult to perform, and could affect obligations that way. Second, other agents not doing their 'fair share', will often make the *best* action an agent could perform considerably more demanding, and if we think morality cannot demand more than what is fair (in some sense), theories which do not take this into account will be inadequate. Third, we might think that the social factors have some role in *constituting* what our obligations are. I consider these options in turn.

How can the number of people performing some act make it more difficult? And how could this affect what one is actually required to do? We might consider the effect of peer pressure. If many people in a given social group or community perform a certain charitable act – and the act is such that knowledge of those performing it is publicly available or readily communicated – an agent is more likely to feel compelled that they should also perform that act. Not only will good behaviour affect the sense of obligation and willingness to perform such acts, but it also seems like this type of factor may alter the perception of how *difficult* it is to perform.

For an example of this sort of behaviour and the perceptions of it by the agents concerned, we may look at the villagers of Le Chambon during World War II. Lawrence Blum gives this example when discussing the role of communities in sustaining virtuous behaviour (1998).<sup>17</sup> Le Chambon is a French village, which during World War II sheltered around 3,000 refugees. Sheltering refugees was a serious crime at the time, punishable by impris-

---

<sup>17</sup>Blum notes that this is an oft-cited example, but himself draws heavily from discussions of the community by Alasdair MacIntyre (1984, p.10).

onment and potentially death, though none of the villagers of Le Chambon were executed. Three village leaders were jailed during this time for aiding refugees.

It may be surprising that so many villagers helped out with assisting refugees. The effort clearly involved large numbers of the community, as the village only had a population of about 3,000. What is more pertinent for my purposes here is how the villagers later described their actions. Blum notes that though we see the actions of these villagers as extremely admirable, “The Chambonnais did not see their actions and practises as having been virtuous in the ‘noteworthy’ sense. When pressed to say why they helped the refugees, they said, ‘It was simply what one had to do’, ‘She [a refugee] was standing at my door; how could I fail to help?’ and things of that sort” (1998, p.242).

Clearly, however, their behaviour *was* noteworthy. The risks taken by the inhabitants of this small village for the good of strangers were truly exceptional. Given that this type of collective action was not replicated in other similar villages, it seems plausible that the social environment within the village had a huge impact on what was perceived as required of them.

Being surrounded by people acting virtuously can be seen as promoting virtuous behaviour in several ways. The presence of exemplars demonstrates that the virtuous behaviour is possible.<sup>18</sup> When a significant proportion of a community acts in accordance with some virtue, this is bound to be even more pronounced. On the other hand, if a virtuous person is seen as the exception, thoughts like “I could never do that” or “That person is just extremely impressive” might make it more plausible for an agent to believe they simply are not able to act so virtuously. If surrounded by people similar in all salient factors to oneself, it seems apparent to a person that *they too* can act that way. Others performing a virtuous act makes it seem more like *we could do it too*. That the presence of others who act similarly makes it easier for us is supported by Hallie’s suggestion that we “fail to understand what happened in Le Chambon in we think that *for them* their actions were complex and difficult” (1979, p.284).<sup>19</sup>

---

<sup>18</sup>This thought is noted and substantially developed in exemplarist virtue theory, e.g. Zagzebski (2017).

<sup>19</sup>Vanessa Carbonell argues that the presence of moral saints can increase our moral obligations. Witnessing others perform saintly acts, bearing heavy burdens or making huge personal sacrifices, provides us with “a special kind of evidence that bears on what

When we consider acts that are available to us, it seems very plausible that how difficult it is to *commit* to an act depends on how difficult one *regards* the act. If then, the virtuousness of others in our local environment alters our perception of how difficult an act (or lifestyle) is, it also alters how difficult it is for us to actually commit to that act. It is natural to think it is more *demanding* to perform hard actions than easier actions. Accepting this, it seems that the easier a given good action is, the lower the benefit would have to be to make it morally required. We can imagine reacting particularly harshly to an agent who chooses not to help someone in dire need when doing so would be very easy, and doing so because they could *so easily* have helped. In this way, being surrounded by a virtuous culture (at least indirectly) alters what one's obligations are.

We may also look at how an agent's obligations might change from a lack of virtues within a community. While in the case of Le Chambon, the general community may have found it easier to act virtuously, if an agent is surrounded by others who seem to do little or nothing to help others, they may respond to feelings that they should help by comparing to others. Though recognising that it would be good to act in an other-regarding way, they may think "No one else is helping" or "I'm doing as least as much as those guys". Seeing no exemplars, they may perceive the acts in question as so exceptional that acting like that is something they just couldn't do.

Is this type of difficulty accommodated by existing moral theories? It seems that the only way ordinary (welfarist) consequentialist theories may count this difficulty as a reason 'against' performing an action is if it is regarded as some cost to wellbeing. However, it is not clear that this can be done. It seems extremely implausible – or at least I think so – to view the difficulty of an action one has performed, *ceteris paribus*, as affecting one's current welfare.<sup>20</sup>

Another way social factors could (conceivably) affect one's obligations arises when we consider fairness within a community. Someone may agree with Singer in thinking that everyone in the western world should donate a proportion of their disposable income so that collectively we might eliminate

---

we can reasonably believe about our obligations" (2012, p.228). While Carbonell does not put her thoughts in terms of difficulty, instead viewing exposure to moral saints as removing ignorance, much of what she says seems expressible in the terms I have used here.

<sup>20</sup>This question is addressed further in chapter 5.

famine, or deaths from certain diseases. *But*, they might think, very few others in the western community donate *that* amount, which means that to accomplish the same good, those who are donating would have to give a lot more. They may wonder “Why should I give away so much of my money, just because other people aren’t doing their bit?”

The types of thought above suggest that there are some ways that the expectations or compliance of others might affect one’s obligations. However, this might strike us as puzzling. Why should other people acting morally or immorally have an effect on how demanding my obligations are?

Singer discusses this in “Famine, Affluence and Morality”. He compares his basic pond case with the global poverty case and notes that there is an asymmetry, because in the pond case there is only one person who can save the child, whereas there are millions of affluent westerners doing nothing (or negligible amounts) to aid the global poor. When discussing whether this asymmetry is morally salient, he discusses an amended pond example, where there are several onlookers doing nothing. As it seems we would be under no less of an obligation in that case, Singer concludes that we should see “the absurdity of the view that numbers lessen the obligation” (1972, p.233).

However, simply because in this particular situation, the numbers of potential saviours does not seem to affect what one must do, this does not mean fairness does not have some role in determining what one ought to do. Singer’s amended pond case does not obviously lend itself to a sensible division of labour, whereas in the case of global poverty, we can divide the costs of saving people between those who are able and arrive at a fair amount (perhaps a percentage of income, or of disposable income). In the pond case, it does not make sense to see the saving of the child as a group project. If, as Murphy (2000, p.96) and Cullity (2004, p.78) suggest, this is an important distinction, Singer’s ruling out of numbers having a role therefore seems too quick.

If an obligation-generating theory takes no account of social factors which affect agents’ perceptions of obligation, they are likely to miscalculate what obligations agents actually are subject to. Murphy provides a response to demandingness objections which at least attempts to attempt to account for some social factors. His particular response will be considered in 4.2.

I have suggested in this section that factors of local expectation and

compliance can affect an agent's obligations in a variety of ways. The presence (or absence) of exemplars may reduce the perceived difficulty of an action, consequently reducing its *actual* difficulty, which in turn may affect its obligation-status. Behaviours being seen as common or expected within a community may also make those actions comparatively easier because of threats of social sanctions if one does not 'toe the line'. That is, not acting in accordance with some norm carries a risk of drawing criticism, and complying with that norm will not carry that risk.<sup>21</sup> It also seems that there is something *unfair* about a requirement to 'pick up the moral slack' when other agents fail to do their bit, though it is not obvious if or how a moral theory should respond to this.

If a moral theory does not factor any of these social features into its assessment of which acts are obligatory, as I've suggested that typical welfarist consequentialism will struggle to do, it may be impoverished as a result. Given the variety of ways considered here that social factors might affect our obligations, and that these are factors which seem unaccounted for by standard moral theories, it could be that a failure to appreciate these factors will, in any otherwise plausible theory, make it overly demanding in some instances.

It is difficult to evaluate whether this type of inadequate sensitivity to social factors is necessary to make a theory overly demanding, simply because it is unclear what a proper appreciation for these factors would entail in a theory. As with previous considerations, we might be tempted to conclude that an insensitivity to social factors cannot be necessary for demandingness on the basis of iterative cases. We might think that as long as we commit to the claim that 'an agent has an obligation to perform an action involving negligible cost to relieve significant suffering' then by the parity of similar cases, which may be iterated, an extremely demanding conclusion is inevitable. However, if the way an obligation is generated is sensitive to issues of fairness or general compliance, the iteration may be halted, by claiming that an agent has done one's part or contributed significantly.<sup>22</sup>

---

<sup>21</sup>It is notable that if agents believe practices are more common, they are typically more likely to believe that these practices are permissible Lindström et al. (2017).

<sup>22</sup>This type of thought is considered later, particularly in sections 4.2 and 6.4.

### 3.4 Insensitivity to Individual Psychology

As well as social factors that may make a course of action more difficult for an agent, we might also consider aspects of an individual's psychology to have a similar effect. In this category I have in mind particular attachments, character traits, fears or predilections, which might act as barriers to performing an action one otherwise would seem obligated to perform. It seems to be a natural thought that these features do affect what we consider to be required of agents in certain situations.

In this vein, Richard Miller argues that people should be giving more than they currently do, but that this should not require sacrificing attachments to goals that enrich one's life (2004, p.361). As long as the projects are "worthwhile goals", Miller suggests that a moral theory should not demand their sacrifice (2004, p.362). Because of this, people psychologically constituted in such a way that their lives would be made significantly worse very easily – perhaps their worthwhile goals require practically all of their available resources – would (according to Miller) have lesser obligations than other people.

While this is controversial, there is a plausibility to the notion that having a certain attachment to a goal or activity might count against actions that would be incompatible with their enjoyment. To take a more modest example, consider Daphne, an agent who is usually quite sensitive to the plight of the world's poor. Daphne gives a large portion of her income to charity, but has one family heirloom which she would not sell, despite knowing that it would raise enough money to do a great deal of good. She would find it hard to part with the item because of its sentimental value to her. We probably would not begrudge her this one luxury that has strong personal value, despite her knowing she could save several lives if she did part with it. If we find examples like this convincing, it suggests that psychological attachments can grant 'special permissions' to not perform some act that would otherwise be obligatory.

Similarly, it seems plausible that people who are generally of a friendly disposition but find certain beneficent acts – like helping with the elderly, perhaps – particularly difficult might have less of an obligation to perform

those actions.<sup>23</sup>

Another type of example we might consider could involve confronting phobias. Imagine Aga, who is mildly agoraphobic. She is uncomfortable in open spaces, though the anxiety subsides very quickly once the situation is over. One day, spending her time in her house conducting her standard daily routine (all indoors), she notices through the window that a baby is in the middle of the road. No one else seems to be around, so she knows that to secure the baby's safety she would have to go outside for a brief time. The very thought of doing so makes her panic, but she knows she could do it.

Though we might well think Aga is still required to rush out to save the baby, her condition definitely seems to complicate matters. If she did not go outside, maybe we would blame her, but not to the same degree as a neighbour across the road without such a condition who watched idly.

An alternative way of accounting for this phenomena is to treat these types of psychological features as excuse conditions.<sup>24</sup> The 'right' actions in each case were still obligatory, but the agent is excused (or their culpability is mitigated) by the extenuating psychological circumstances. A potential failing of analysing these situations as only excuse conditions is that we might be tempted to see the agents as no more praiseworthy (than a person without the specified psychological feature) if they do act in the right way – they have still merely done what is obligatory.

### 3.5 The Diagnosis

The availability of a demandingness objection to a moral theory depends on that theory making demands of agents. This much is clear, and it does not seem to matter what kind of obligations these are; they may be coarse-grained or fine-grained. A plausible theory that also maximises — that requires maximal promotion of whatever goods it espouses — seems at least *likely* to be overly demanding.

While embracing impartiality is often heralded as leading to demanding obligations, there are several ways in which a theory may be rightly labelled 'impartial' that do not have this result. The only notion of impartiality

---

<sup>23</sup>Whether difficulty has an effect on moral obligations is discussed extensively in chapter 5.

<sup>24</sup>This is the solution I favour, and argue for in section 5.2.2.

which would guarantee that a theory is demanding — impartial benevolence as the direct guide to action — is not a feature possessed by many theories, so is clearly not *crucial* in propagating extreme demands.

I have suggested that factors involving one’s moral community, such as local compliance and expectation, can (at least indirectly) change how demanding a situation is. If theories do not accommodate this (and many do not) the verdicts of how much is demanded in those situations is likely to be mistaken.

It is also apparent that our judgments about moral agents do vary in accordance with various features of individual psychology. However, I tentatively suggested that these considerations may be best understood as playing a role in granting excuse-conditions. It certainly is not the case that unusual features of individual psychology explain most cases of overly demanding verdicts, though, as psychologically ‘normal’ agents would (and do) find many of these verdicts implausibly demanding.

In terms of isolating a cause for demanding verdicts, the most promising type feature from this analysis seems to be an insensitivity to certain features of agent’s moral community.

### 3.6 Conclusions

In this chapter I have attempted to find which features make a theory vulnerable to this objection. The most obvious feature is a generation of moral requirements. These requirements might be disjunctive (obligations to do  $\phi$  or  $\psi$  or  $\xi$ ). Though impartiality is often cited as being important, I have argued that a theory need not be impartial in order to fall prey to demandingness, nor will every impartial theory be susceptible.

Instead, a failure to account for several broad factors – social features which might affect obligation and certain aspects of individual psychology – were noted as promising avenues. If what I have suggested is correct, then issues of compliance and expectation within a moral community, and an agent’s perception of these features have a much greater role in determining what an agent’s moral obligations actually are than is usually accepted. Psychological differences among individuals were also discussed, though no verdict was concluded about whether an ignorance of those issues can result in overly demanding accounts. In the following chapters, I consider moral

theories offered in response to overdemandingness which can be seen as taking these features into account.

## 4

# Potential Solutions and Failings

*“Everybody Else is Doing it, so Why Can’t we?”*

– The Cranberries, album title

*“A problem left to itself dries up or goes rotten. But fertilize a problem with a solution – you’ll hatch out dozens.”*

– N. F. Simpson, *A Resounding Tinkle*

In this section I discuss two recent attempts, by Miller and Murphy respectively, to provide responses that adequately address the demandingness objection by making changes in accordance with some of the features mentioned above. Richard Miller attempts to soften moral obligations by examining underlying dispositions rather than individual actions. This move is able to cater for differences in individual psychology leading to different moral obligations. Liam Murphy limits demands in relation to social factors, particularly focussing on how demands change in nonideal theory, i.e. in situations where there is not full compliance with the moral rules.

I will suggest that each of these attempts ultimately fail, and attempt to illustrate why this is the case, before considering what lessons might be learned from these accounts.

## 4.1 Miller's Underlying Dispositions

Richard Miller sees his project as a direct reaction to Peter Singer's extremism. He concedes that while Singer's principle looks extremely plausible, upon reflection, its demandingness "misconstrues ordinary morality" (2004, p.358).

In opposition to Singer, Miller proposes a more moderate account of our duties of beneficence, based on the 'Principle of Sympathy',<sup>1</sup> that:

"One's underlying disposition to respond to neediness as such ought to be sufficiently demanding that giving which would express greater underlying concern would impose a significant risk of worsening one's life, if one fulfilled all further responsibilities; and it need not be any more demanding than this."

(Miller, 2004, p.359)<sup>2</sup>

Rather than focussing on individual acts, Miller's principle examines the appropriate attitude of concern an agent should have. This concern is an "underlying disposition", which should be demanding enough that expressing more concern would risk making one's life worse.<sup>3</sup> Miller is clear that the suggestion that minor reductions in the quality of goods in some instances or other trivial events do not count as worsening one's life. Against the possible counter-claim that such trivialities do worsen one's life, but only by an insignificant amount, he doesn't offer a rebuttal, but instead suggests that such opponents recalibrate the view to treat "significant risk of worsening one's life" as short for "significant risk of significantly worsening one's life" (2004, p.359).

Miller makes a distinction between underlying dispositions, which express basic concerns in a coarse-grained way, from personal policies, which

---

<sup>1</sup>Henceforth *Sympathy*

<sup>2</sup>Miller also defends this principle in the first chapter of *Globalizing Justice: The Ethics of Poverty and Power* (2010).

<sup>3</sup>This may seem reminiscent of Kant's imperfect duty to beneficence. Kant holds that "beneficence is the maxim of making others' happiness one's end, and the duty to it consists in the subject's being constrained by his reason to adopt this maxim as universal law" (*The Doctrine of Virtue*, in *The Metaphysics of Morals*: §9). Imperfect duties, such as this, involve latitude, so may be discharged in various ways. Though there is some disagreement among Kantians (see Timmermann 2005), several argue that for Kant, one is not morally required to help in every situation (e.g. Baron 1987, Hill 2002).

are “specific standing commitments” (2004, p.360). Underlying dispositions are the subject of *Sympathy*. A similar disposition may, in different people provoke a different response, based on one’s inclinations, means and skill set. We may consider the disposition to help those in need in a situation where a bystander has a medical emergency. The underlying disposition will prompt one to help. If the would-be helper has medical training, this help will probably be manifested by immediate medical help, whereas if not, calling for help might be the more appropriate personal policy.

The more moderate demands entailed by the *Sympathy* depend on what counts as significantly worsening one’s life. Miller wants *Sympathy* to be more compatible with everyday moral commitments, thus allow agents to act in accordance with various worthwhile goals the agent identifies with, “with which he could not readily detach” (2004, p.360). For Miller, abandoning a worthwhile goal, such as developing a sophisticated appreciation of music or expanding one’s culinary horizons, does significantly worsen one’s life. Clearly, *Sacrifice* would not allow agents such indulgences, but Miller’s *Sympathy* does encourage people to pursue “worthwhile goals”. These goals are seen as basic interests; those that give “point and value to specific choices and plans” (2004, p.361). Miller’s “basic interests” must not be confused with the interests an agent needs to merely survive, but are those interests that a person may see as essential to who they are.

Miller even defends some extravagant interests, even when their pursuit comes at the cost of donating more to the needy. Though it would allow us to relieve suffering further if we had more modest interests, once one has certain interests and one’s life would be worsened if they cannot be pursued, one should not be obligated to damage one’s life by abandoning them. This allows *Sympathy* to be compatible with the occasional (and depending on the nature of the interest, even the not-so-occasional) purchase of luxury goods.

Despite this relaxing of obligations, Miller is adamant that *Sympathy* would require significant spending from most of the world’s non-poor, though he does acknowledge that some people with expensive tastes (e.g. if one has an “exorbitantly expensive goal that most of us lack or could readily do without”) will need to “retain much more than he otherwise would to avoid worsening” their lives (2004, p.363).

Miller is even able to defend why this particular principle is appropriate.

He sees equal respect as the foundation for judgements of moral wrongness: “a choice is morally wrong if and only if it could not be made under the circumstances by someone displaying equal respect for all persons” (2004, p.366). If equal respect for all is a good foundation from which moral claims can be deduced, and *Sympathy* is consistent with equal respect for all – that is, it does not provide contradictory verdicts – then *Sympathy* must be a good moral principle (at least with regards to what is morally wrong).

One might wonder how it is possible to respect everyone equally, yet allow some in the third world to starve, and privilege others, like oneself or one’s family. To alleviate this concern, Miller notes the distinction between equal respect and equal concern. Equal concern would require that one be equally distraught by a stranger contracting a deadly illness as her child catching the same illness (which strikes most of us as undesirable). Though we may respect everyone equally, the interests of those closest to us will bring us to act more readily and more acutely than those of a stranger. Similarly, we will (generally) be more readily moved by our own needs than by those of others: Respect for others need not require us to “be prepared to do violence to who one is” (2004, p.362). As we identify with our worthwhile goals – we see our goals as part of who we are – and we should also respect ourselves, we cannot be required to abandon these worthwhile goals. So, Miller is able to argue, equal respect for all is consistent with *Sympathy*.

Miller is adamant that equal respect is universally acknowledged as good, and explains that it does not entail equal concern or equal treatment, but it isn’t exactly clear what this lofty notion – ‘equal respect’ – actually means. It seems that Miller sees this as suggesting to us all a particular way of thinking about agents and requiring a degree of consideration. The term itself, “equal respect”, appears compatible with not respecting anyone at all, as long as everyone is not respected equally.

It might be that equal respect is a confused notion, but I will suppose it must involve certain features. Primarily, it must involve some appreciation that there is no objective reason to favour the interests of one person over another. It would be inconsistent with equal respect to promote government policies belying a valuing of a certain culture, person or group as more valuable than another. When one manifests unequal concern, this must be because of some special relationship, and not because of a supposed difference in intrinsic value. Equal respect must also entail an acknowledgement

that all moral agents have interests of their own, and a recognition that their goals have value for them in the same way our goals have value to us. I will assume for the remainder of this paper that “equal respect” is able to perform the justificatory function he supposes, though I do acknowledge that further probing could be undertaken in this avenue. Equal respect for all, we may suggest, is a common-sense moral principle that we all do and should support.

It is perhaps of interest to note that *Sympathy* does not generate any of the counterintuitive judgements to several of the counterexamples posited against *Sacrifice* discussed in 2.2.1. One such type of example involves iterative demands. If one can save a child from drowning at negligible cost, it seems obvious that one should do it. However, even when the stakes are this high, if there is an endless amount of children dying, we might think that after a certain point, it is permissible to not save one. This point has recently been argued by Travis Timmerman (2015). According to Miller’s account, what would be morally required would be having an appropriate level of underlying concern. This would require a serious and genuine concern, but would allow one to give up at some point, acknowledging that manifesting a degree of concern more than this would significantly worsen one’s life.

#### **4.1.1 Problems of Miller’s Account**

There are several problems with *Sympathy*, some perhaps providing only minor difficulties that could be reconciled. I argue here that some of these problems require us to deny *Sympathy*. Though *Sympathy* is supposed to provide us with a good moral principle to live by, it seems too generous to some (and possibly not to people we deem morally good), too harsh on others and, after some considerations, unclear.

##### ***Sympathy* seems too weak**

Because of Miller’s stance on “basic interests”, *Sympathy* may make very minimal (or no) demands of certain agents. A basic interest gives “point and value to specific choices and plans” and there is no restriction on how costly or affordable these may be (2004, p.361). The basic interests, and the goals an individual has as a result, could be quite ostentatious. If one of my goals involves expensive travel, my life will be made worse by my inability to

do so, so I could choose to opt to engage in very expensive trips around the world, rather than giving to the poor. We might imagine that this expensive travel is of such a high level that no disposable income remains.

It might be suggested that this example is uncharitable. Maybe this sort of ostentatious travel or other extremely expensive activities should be excluded as possibilities for worthwhile goals that contribute to one's basic interests. Miller clearly envisages *Sympathy* as usually justifying only "occasionally purchasing...luxuries and frills" (2004, p.361), so it might be supposed that some limitation on what could suffice would be beneficial. For several people with such hobbies, Miller would claim that while such activities are enjoyed, their loss would not constitute an actual worsening of one's life. It is important not to overrate "what merely frustrates, blowing it up into something that worsens one's life" (2004, p.362). Several people may profess that because of their deep love of extravagant travel, its sacrifice would harm their lives substantially, but in doing so may merely be overstating or overestimating its importance to them.

Though some people, after reflection, may concede that they had overrated the value of ostentatious goods to their lives, it seems that some certainly would not. Some people *really will* identify with exorbitant habits, and see their lives as thwarted if they needed to abandon them. We might imagine someone who has the goal of visiting every capital city in the world, and exploring there, as a means to enrich herself. If she was to contemplate missing out on her extended expensive trip around South East Asia, she might feel a great sense of loss. So, rather than donating tens of thousands of dollars to an effective charity and potentially treating over 5,000 people for neglected tropical illnesses for a year,<sup>4</sup> *Sympathy* seems to allow her to indulge in her expensive trip.

There seems a difficulty here, as those who have had the means to develop extremely expensive habits throughout their lives must be those who have had the greatest means available. Due to their wealth, these are the types of people who we *really want* to give most to the world's needy. However, these are also the people who will be allowed to spend the most on ostentatious goods under *Sympathy*, as understood thus far. It seems like *Sympathy* could be supplemented by additional premises governing which goals are

---

<sup>4</sup>Based on a donation of \$10,000 to the Schistosomiasis Control Initiative (figures from [www.givingwhatwecan.org](http://www.givingwhatwecan.org) in July 2017).

worthwhile, on the basis of just how costly they are. Surely, those who spend millions on expensive art collections – when there are vast communities suffering acutely who could be relieved of their suffering with that money – are acting in an unjustifiable way? Even if we acknowledge that their donating would risk making their life worse, given the importance of what is at stake, they ought to do that.

However, Miller does argue quite explicitly against this when considering characters like Richard Dworkin’s ‘Louis’. What Louis cares about is dominated by “an extremely expensive, worthwhile goal of ultra-refined savoring of food and drink meriting such sensory discrimination” (2004, p.363). As his goal is worthwhile, the demands of *Sympathy* will be minimal, or none. Despite many of his neighbours living moderately austerely, Louis can revel in his expensive tastes. Nobody is required to aid Louis in his pursuit of these goals,<sup>5</sup> as that certainly would not be required by our responsiveness to the neediness of others, but he must be allowed to use all the resources at his disposal to fulfil his goals.

Perhaps, though we may judge this as unfair, Miller would simply concede that such people would not be obligated to give as most people are. Sacrificing these particular life goals could be seen as doing damage to who Louis is as a person, and if he was to give to the needy it would be at the cost of his own identity.

This type of response, however, does seem untenable. We may imagine our other Louis-type character, who has the worthwhile goal of travelling the world. She does so in such a way that commits her to giving less to the poor than someone raised in less wealthy conditions (and who, as a consequence, has less extravagant goals). We can, by stipulation, increase the wealth of the agent and the costs required by her project, such that she is spending millions, in which case it definitely seems like she is worthy of criticism/scrutiny. In these cases, the principle seems too weak. As Elizabeth Pybus notes, “selfishness doesn’t relieve us of the moral requirement to be willing to do things for other people” (1982). Engaging in ostentatious

---

<sup>5</sup>Miller stipulates that nobody has this obligation because the “neediness” involved in *Sympathy* is seen as on a par with “the sort of deprivation that Singer labels “very bad””, i.e. those needing relief from starvation or disease (2004, p.359). Louis clearly does not fit into this category. Exactly why it is only the Singer variety of neediness that *Sympathy* should apply to is unstated, and perhaps warrants its own discussion, but that may prove tangential to current purposes.

projects, even ones that are “worthwhile”, does seem selfish and unable to justify failing to ever act beneficently. This at least seems problematic for Miller’s account, if not a knock-down objection.

### ***Sympathy* is not a clear principle**

Though Miller may simply bite the bullet in the types of cases described above, the following problem demands a response. *Sympathy* holds that our underlying dispositions to respond to neediness should not be more demanding than would “impose a significant risk of worsening one’s life”. Miller does not elaborate what he considers a significant “worsening” of one’s life to be.

For a “worsening” to occur, there must be a relation between two situations, one judged as better and the other worse. To undergo a “worsening” because of a certain choice, the agent must then be seen as selecting the worse of two options. I propose three interpretations for how this may be understood:

- a a worsening of one’s foreseen expected future,
- b a worsening in comparison to everyone else, and
- c a worsening compared to one’s situation before one gives.

#### **(a) Worsening of a foreseen future**

This interpretation relies on some notion of an “expected life” which an agent’s life could be better/worse than. This particular version seems to have most support from Miller’s comments: When introducing the principle, Miller claims that the significant risk of worsening one’s life means a “nontrivial chance that one’s life as a whole will be worse than it would otherwise be” (2004, p.359). Even this option seems slightly ambiguous. If we suppose we are considering donating a certain amount of money to charity, is the quality of our life under this level of giving supposed to be compared to giving nothing, or to our current level of giving (supposing that we are already donating some amount)?

The first of these options seems likely to be problematic when evaluated by people on middle-incomes who give a small to moderate amount already. Consider Noah, who earns a modest income – enough to survive, put some

money away for a rainy day and have a little disposable income to spend on the occasional luxury – and through his life has been moved by many of the disaster appeals or OXFAM adverts and gives 10% of his income. He also has a passion for nice whisky. Upon thinking about *Sympathy*, he concludes (rightly) that his enjoyment of whisky is a worthwhile goal. He also calculates that if he ceased to give at his current level, he would be able to buy an additional nice bottle of whisky every month, with the hope of expanding his pallet.

Most people will seem to have an interest or passion similar to Noah's, one that they will identify with and that will enrich their lives if properly pursued. It might also be the case (and I suspect it is)<sup>6</sup> that the vast majority of these could pursue such goals much more successfully (to the extent that lacking this would make them substantially worse off). If correct, this would lead not, as Miller claims, to people giving much more of their incomes to the needy, but potentially much less.

The second suggestion for this interpretation might result in situations more similar to those Miller envisages. Noah, and the vast middle-classes, would judge whether giving an additional amount to their current donations would lead to the pursuit of their goals being substantially less successful and consequently worsen their lives. However, basing what should be given merely on one's current giving practices seems somewhat arbitrary, and Miller claims that this is something he aims to avoid (2004, p.370). This also might seem to unfairly disadvantage those who are currently donating a non-negligible amount of their income, as they would not be comparing to an egoistically ideal possible future.

### **(b) Worsening of a future in comparison to others**

One possible interpretation that might be taken could set the ordinary, status quo of giving as the grounds for comparison. This interpretation seems the least plausible, but still worthy of mention. It might be thought that everyone in the west should be donating, but everyone in a similar position should be donating a similar amount. If a small amount of people donate substantial amounts of their earnings, while others earning similar amounts in their community give little or nothing, those who do donate may find

---

<sup>6</sup>This is an empirical matter, and also subject to rationalising, which may add difficulties to the application of *Sympathy*.

themselves in relative poverty. We might see that as objectionable, so conclude that one's donations shouldn't leave them worse off within their own local society.

This interpretation seems to allow for a "collusion" of not giving, so that no one worsens their lives compared to everyone else in their position. As Miller's principle is also couched very cautiously – even *risking* worsening one's life is motivation enough to warrant not giving additional money – it might even lead to people reducing their donations, because otherwise they may risk worsening their position in society.

### **(c) Worsening in comparison to pre-giving state**

One natural interpretation of the "worsening" Miller refers to is in comparison to one's state before they make their donation. As I have already mentioned, Miller does suggest that the comparison should be against what an agent's situation would be otherwise, not before. However, this interpretation should be mentioned, as it would be an alternative if it did not pose significant issues of its own (though I will suggest that it does).

Independently of any donations to charities, people will experience increases or decreases in their standards of well-being. Whether due to changes in income, the cost of living or any number of non-financial factors, an agent's well-being is unlikely to be constant. People in relative poverty may become quite affluent, and vice versa. On one extreme, we may imagine a young socialite from a very wealthy family, but who earns nothing. Despite inheriting several houses, her life is very likely to be substantially worsened if she gives money to charities, simply because her quality of life is likely to worsen anyway, certainly over a long period of time. Perhaps several of the family homes will have to be sold to pay for a very expensive lifestyle (full of worthwhile goals, we may assume for the sake of this argument) and this is certain to impact her quality of life.

This interpretation may then lead to the unsatisfactory conclusion of some of the wealthy elite having no duties of beneficence. The example could also be ran the other way. A moral agent trying to improve her station in life would be obligated to give at a level such that it wouldn't risk worsening her life from that starting position. If an agent's life is unlikely to be worsened – because of a poor starting position – no matter how hard she works, this interpretation would require her to remain in a similar position.

As well as obviously being contrary to Miller's purpose, a principle like this would impede social mobility to a grotesque degree.

### ***Sympathy* seems too demanding**

Given that the first objection I raised to *Sympathy* was that it was too weak, it might seem peculiar to claim that it is also too demanding. However, just as the principle seemed to weak with regards to a certain group of society (namely, rich people with ostentatious hobbies), it is (at least sometimes) too harsh on less wealthy people.

As mentioned previously, it seems clear that Miller did intend *Sympathy* to require agents to compare how their life would be if they donated extra to a possible future that could be expected if they did not. This might also be the most plausible interpretation of the principle. However, despite the earlier objection that *Sympathy* might be too weak, this interpretation can actually be very demanding for certain moral agents.

I suggest this because of how we are likely to envisage expected possible futures. Consider Sally, a child from a western family who grows up in relative poverty. Sally's family is heavily reliant on benefits throughout her youth and as a result she is well-fed and educated, but not affluent. She works very hard and is rewarded by occasional promotions. Sally is humble, however, and whenever greeted by a promotion is surprised, and does not expect to receive another. She does not expect to ever earn enough to satisfactorily pursue her worthwhile goals. Because Sally does not think she can pursue these goals to a degree that would further enrich her life, she regards *Sympathy* as commanding that she adopt a very high level of underlying concern, as she does not believe doing so is likely to worsen her life by hindering these goals.

As Sally's tale is intended to illustrate, this interpretation of *Sympathy* can also result in a hindering of social mobility. We might suppose that if Sally had an accurate idea of how her circumstances would change, she would have realised that she would comfortably have been able pursue her goals. Miller could respond to this criticism by claiming that rather than what an agent actually expects, the principle concerns rational expectation, or what one should reasonably expect. This could mitigate problems caused by Sally systematically underestimating the prosperity in her future. However, to bring out the problem, we need not make use of a case with an

agent who underestimates their potential to fulfil their goals as much as Sally does. If an agent becomes accustomed to a more austere lifestyle, they are unlikely to identify with more glamorous goals, even if they may be attainable.<sup>7</sup> Childhood dreams of pursuing an exciting career are likely to be left unfulfilled by most. An alternative take one could take on *Sympathy* might suggest that even childhood dreams or fantasies do count as worthwhile goals one identifies with, but this interpretation would minimise to the point of negligibility any duties of beneficence in such cases.

#### 4.1.2 Rejecting Miller's view

Though *Sympathy* is able to respond to several of the problems levied against *Sacrifice*, it carries a lot of baggage of its own. Even when the principle is properly specified, the obligations it generates seem extremely unjust. It strikes many as obvious that the bulk of our obligations to the poor should fall at the hands of the rich, as they are the most able to do something about it. *Sympathy*, however, seems to many of the rich 'off the hook', while having the unintended consequence of restricting the spending of the poor to the minimal luxuries they may have become accustomed to.

The switch from Singer's principle, focussing on individual actions, to a principle which concerns underlying dispositions, appeared to be promising. Because of this, it is able to satisfactorily avoid several of the problems *Sacrifice* faces and give appropriate responses to various problem cases.

One of the purported successes of this type of account is that it can give different verdicts of what an agent's obligations are due to individual psychological differences. In section 3.4, whether factors like this – including things like phobias and special attachments – could make a difference to one's obligations. However, given some of the examples which appear problematic for Miller, it does not look like this is promising – at least in the way that Miller attempts to account for the obligations.

---

<sup>7</sup>We might consider Isaiah Berlin's example of an agent who has adapted their preferences to be so minimal that their desires are all satisfied even under a tyrant (Berlin, 2002, p.182).

## 4.2 Murphy's Collective Beneficence

Liam Murphy offers an account of our moral demands that takes into account some of the factors I termed “social factors” in the previous chapter. He observes that the obligations prescribed by many demanding moral theories are, in many situations, extremely demanding because agents cannot expect others to comply with the theories. Because of this, agents who do attempt to comply with the theories are subject to unfair demands. In this section, I describe how Murphy arrives at his proposed principle, which he cautiously offers as “more plausible than the principles” he is aware of (2000, p.5), before discussing some undesirable implications of the account and consequently rejecting it.

In Murphy's discussion, he focuses on issues pertaining to beneficence. It should be noted that duties of beneficence are not the only way extreme demands are generated by moral theories, so his response might not be available for all potentially problematic cases. He illustrates the problem by considering an obvious candidate for a principle governing our duties of beneficence, which he swiftly rejects. The ‘Optimising Principle of Beneficence’...

“...requires agents to do the best they can for others. It requires us to keep benefiting others until the point where further efforts would burden us as much as they would help the others.”<sup>89</sup>

(2000, p.6)

This would clearly be, given the current state of the world, an extremely demanding principle. In fact, it looks like it would be even more demanding than Singer's *Sacrifice* Principle, because one would be under an obligation to aid until the burden imposed is *as much* as others would be helped, whereas Singer only requires that the additional sacrifice be “comparable”.

---

<sup>8</sup>This principle is characteristic not only of utilitarianism. Murphy notes that it is also included in the pluralisms of W. D. Ross's, Samuel Clarke and William Godwin (2000, p.10,n.4).

<sup>9</sup>Elsewhere, Murphy offers (and rejects for the same reasons) a similar principle, which he terms the ‘Simple Principle of Beneficence’. This states that every agent is required “to perform the action, of those available to her, that will make the outcome best” (1993, p.268). This is susceptible to all the problems of the optimising principle, but is committed to a consequentialist account of goodness.

Intuitively, Murphy notes, people think this ‘just couldn’t be right’, that the obligations imposed by it are far too excessive (2000, p.6). However, if *everyone* or nearly everyone did comply with the Optimising Principle (a situation he calls ‘optimal full compliance’) the demands on each of us would typically be fairly minimal. The reason that the Optimising Principle of beneficence *is* so demanding is that others are doing much less than they would be required even under full compliance (2000, p.13).

Accepting that the Optimising Principle demands too much of agents (in these cases of partial compliance) is Murphy’s starting point. This gives us reason to reject that principle and seek a new way of understanding morality’s demands. One way that the extreme demands entailed by the Optimising Principle could be curtailed is by placing a limit on what might be demanded.<sup>10</sup> Murphy discusses a class of principles – “limited principles of beneficence” – which do just that.

The simplest way a limited principle might function would be to set some limit to demands, such that no situation could ever require more of an agent than this limit. If this was the only limiting factor, however, it would be inadequate for our purposes. To see this, we can consider instances where it seems that morality really is extremely demanding, like Ashford’s machete case – where “the only way of stopping billions of people suffering an agonizing death was by hacking off your left leg with a fairly blunt machete” (Ashford, 2003, p.274). If we accept that this demand is legitimate,<sup>11</sup> then the limit could not be lower than that amount. But then this limit seems of no help, because even the Optimising Principle *usually* demands nothing like this much of us.

Murphy considers a more complicated type of limiting principle via Scheffler’s agent centered prerogatives. These (by themselves) are not limiting principles because they only provide a multiplying factor by which

---

<sup>10</sup>How this could be measured is a difficult question, and one Murphy spends a significant portion of *Moral Demands in Nonideal Theory* considering. His answer relies upon comparing how one’s life would go under full compliance with some principle from now on, with how it would go if one lived perfectly prudentially in the actual world. The *demand* of the principle is taken to be how much better off one would be in the prudential situation than in the case of full optimal compliance (2000, p.42). It should be noted that this way of viewing demands might have some very peculiar consequences – possibly suggesting that even the Optimising Principle makes no demands of us – as pointed out by Tim Mulgan (2003, p.115). I will not engage with this difficulty here.

<sup>11</sup>If we do not share that intuition, we could give another situation where an enormous amount is at stake where some incredible sacrifice really is demanded of an agent.

one can weight one's own interests when determining what is permissible. However, Murphy suggests that we might amend this view into a limited principle. To do this, the multiplying factor an agent is able to apply to one's own interests could be indexed according to how much good there is to be done. To clarify, in some situations where significantly more good could be done – after an environmental disaster, for instance – the extent one is permitted to prioritise one's own interests would go up, in order to keep the demand below a certain limit (2000, p.65). Even if this were done, however, Murphy suggests that an additional indexing would be required to compensate for the amount of non-complying agents. This double-indexing would not only be a strange manoeuvre, but it would also result in the account failing to capture one of the intuitive benefits of Scheffler's account – that the amount that might be required of an agent increases depending on how much good one can do.

These two approaches – setting an absolute limit and the more complex index-based approach – illustrate a tension that limited principles face. We both want to recognise that it is possible to ask *too much*, yet we also think that our obligations are sensitive to the amount of good that can be done.

In Murphy's view, it is not actually the extreme demands that are the problem with the Optimising Principle, at least not directly. One motivation we might have for thinking this is right is the multitude of legitimate ways we can be subject to extreme demands. For example, we probably think you can legitimately take on huge moral burdens not related to beneficence (e.g. debts, reparations, obligations from promises, parenthood). For reasons like this, Murphy rejects the notion that there is a fixed *limit* to moral demands (2000, p.15). Rather, he thinks the problem is that an ideal theory will involve an “unfair imposition of responsibility” (2000, p.99).<sup>12</sup>

Murphy suggests that duties of beneficence are best regarded — and naturally regarded — as group obligations. He notes that agent-neutral

---

<sup>12</sup>Michael Ridge, drawing upon Murphy's arguments, also suggests that it is *unfairness* that is the problem really identified in moral theories perceived as overly demanding. He rejects Murphy's solution (2009, p.196) and develops his own alternative, which intends to distribute burdens *under partial compliance* fairly, in what constitutes a “kind of compromise between unconstrained beneficence (Singer and Unger) on the one hand, and beneficence as strictly constrained in conditions of partial compliance (Murphy) on the other” (2009, p.198). I do not discuss this here, but I deem it susceptible to several of the criticisms I make in the section that follows, most obviously (though not only) to instances of overdemanding verdicts upon single agents.

principles can be seen to apply to people as individuals *or* as a group (2000, p.96). Agent-neutral principles are those with aims that can be specified without referring to the agent. These are contrasted with agent-relative principles, which must apply to individuals, such as those commanding promise-keeping. An agent who accepts that promise-keeping is important has the goal of not breaking *their* promises, not merely minimising the amount of promises that are broken.<sup>13</sup>

If duties of beneficence are seen as an individual aim, Murphy suggests, the compliance of others can be seen as affecting one's obligations the same way as natural phenomena might. With non-group aims, we respond to what others are doing, but it has no special moral relevance. Just as a promise I make (an individual aim) to make an appointment may become more burdensome if the weather gets considerably worse, if I hold the aim of poverty relief as an *individual* aim, it is irrelevant whether others are complying or non-complying – what matters is how much I can do. Viewed this way, an increase of the number people suffering from an illness would increase my burden the same way a reduction in compliance would.

If however, we conceive of the aim as a group aim, it makes sense to divide the burden among ourselves. When we view principles of beneficence in this way, it makes sense for us to care about how we divide the moral labour. This means that, unlike with individual ventures, it makes sense to ask what we ought to do in cases of partial compliance. When we consider group aims, our “concern that the compliance effects of pursuing the principle be fairly distributed is parasitic on that prior moral aim” (2000, p.106).

Murphy thinks this is how we should view our duties of beneficence. One reason that seems to count in favour of this way of viewing beneficence is the agent-neutral character of the obligations. If we really have the aim of eliminating starvation, we do not care *who* feeds the starving. It should not matter to an agent “if the benefits flow more through the agency of others than her own agency” (2000, p.75). We could find a related motivation in the way certain members of the public think that duties of aid are ‘for the government’, suggesting that we do view it in terms of a collective responsibility (2000, p.4). While people may express this type of thought about aid, this makes an obvious contrast with, say the obligation not to kill, where no one would make such a claim.

---

<sup>13</sup>Murphy notes that this example is borrowed from Nagel (1986, p.177).

Another motivation for viewing duties of beneficence this way is that he thinks it is the key to dissolving demandingness objections. I will return to this point shortly.

Because he sees something special about cooperative aims – such as the shared cooperative aim of beneficence – which affect how much is demanded of people under partial compliance, Murphy determines that we require a ‘nonideal theory’. A *nonideal theory* is one that accounts for situations wherein agents either cannot or will not comply with ideal principles.<sup>14</sup> To provide this nonideal element, Murphy suggests that where cooperative aims are concerned, a *compliance condition* is needed in order to mitigate for unfair distribution of demands.

**The Compliance Condition** An agent-neutral moral principle should not increase its demands on agents as expected compliance with the principle by other agents decreases. Demands on an agent under partial compliance should not exceed what they would be (all other aspects of her situation remaining the same) under full compliance from now on.

(2000, p.77)<sup>15</sup>

The general thought behind the condition is that other people not ‘doing their bit’ should not mean that you have to do more. It has some intuitive pull. Murphy considers an agent who asks: “Why should I do more just because others do less? Surely I should only have to do my fair share” (1993, pp.277-278). Kwame Anthony Appiah makes a similar point, claiming that...

“...if so many many people in the world are not doing their share – and they clearly are not – it seems to me I cannot be required to derail my life to take up the slack.”

(2007, p.165)

Garrett Cullity describes views which embrace this thought ‘Fair Share views’. These divide the burden of some collective project fairly among those it falls upon and use this to fix a limit of what is required (2004, p.74).

---

<sup>14</sup>It is this insistence on accounting for factors of expectation and compliance which means Murphy’s account incorporates some of the factors discussed in section 3.3

<sup>15</sup>Also formulated in very similar terms in “The Demands of Beneficence” (1993, p.278).

Any number of moral principles could be devised that would meet the Compliance Condition. For instance, we could propose an extremely minimal conception of morality's demands, which claimed that in full compliance very little (or nothing) is required, and this does not change in partial compliance. Alternatively, and possibly more plausibly, we might suggest a principle which set moderate demands in typical situations of full compliance, but included some limit to the stringency of morality's demands. So long as the principle did not require more of an individual in partial compliance than it would in full compliance, it would meet Murphy's condition.

In accordance with this, he proposes the *Collective Principle of Beneficence*.

**Collective Principle of Beneficence** (simple formulation)<sup>16</sup>: A person need never sacrifice so much that he would end up less well-off than he would be under full compliance from now on, but within that constraint he must do as much good as possible.

(2000, p.117)<sup>17</sup>

The Collective Principle of Beneficence (CPB) would provide the intuitive verdict that our obligations to the global poor in the current state of the world are considerably more moderate than the extremists suggest. It also has the virtue that if some catastrophe (like a natural disaster) were to ensue plummeting additional millions of people into absolute poverty, considerably more would be required of us. This seems in keeping with what we typically think (as is demonstrated by the surges in charitable donations in the wake of major earthquakes or similar disasters).

---

<sup>16</sup>The full formulation, which Murphy admits is 'terribly complex' is as follows:

"Everyone is required to perform one of the actions that, of those available to her, is optimal in respect of expected aggregate weighted wellbeing, except in situations of partial compliance with this principle. In situations of partial compliance, a person's maximum level of required sacrifice is that which will reduce her level of expected well-being to the level it would be, all other aspects of her situation remaining the same, if there were to be full compliance from that point on. Under partial compliance a person is required to perform either an action, of those requiring no more than the maximum level of required sacrifice, that is optimal in respect of expected weighted aggregate well-being, or any other action which is at least as good in respect of expected weighted aggregate well-being.

However, no one is required to act in a way that imposes a loss on some other person unless that other person's level of expected well-being after the loss would be at least as high as it would be, all other aspects of the situation remaining the same, under full compliance from that point on" (2000, pp.117-118).

<sup>17</sup>Elsewhere this is called the "Co-operative Principle" (1993, p.280).

## 4.2.1 Problems with Murphy's View

### Partial Compliance and Small Marginal Sacrifices

The CPB holds that our individual duties of beneficence are derived from group obligations — with the particular requirement upon each individual calculated some fair way — and each agent is never obligated to do more than that.<sup>18</sup> In the case of world poverty, this leads to a verdict which might strike us as plausible.

However, the account does not always provide such intuitive verdicts. In any instance of a group obligation, if everyone in the group complies with a principle, then everyone pays their fair share, and the group solves the problem. If there is only partial compliance — if some do not contribute their fair share, or perhaps contribute at all — the agent is still only required to pay a maximum of that amount.

According to the CPB, the group among whom the collective requirement applies is determined by those who should share the aim. If someone shirks their responsibilities, the obligation is not re-evaluated with regards to the smaller group.<sup>19</sup> Noting this, in his criticism of Murphy's position, Cullity provides an example which he sees as problematic.

“There are three people drowning, and three bystanders, including me. We can rescue them, but only by using a winch mechanism that can be operated by a minimum of two people or, more easily, by three.”

(Cullity, 2004, p.76)<sup>20</sup>

---

<sup>18</sup>It is of course open for debate *how* a “fair share” could be calculated. Do those who are more able to contribute have to pay more – and how much more? Singer discusses a variety of suggestions for how this burden could be shouldered in *The Life You Can Save* (2010, pp.160-169).

<sup>19</sup>If it was, without restriction, the collective principle would seemingly escalate into the optimizing principle for the complying agents.

<sup>20</sup>Recently, Singer and de Lazari-Rabek provide a similar example, where instead of three bystanders and three drowning people, there are ten bystanders and ten drowning children. In their version, you are the only one who jumps in to save a child (2014, p.329). It seems absurd to say that if you save just one and leave the other nine to drown you have satisfied your obligations. This version may be more compelling, but the mechanism leading to the issue – which results in the wildly implausible conclusion – is the same.

In this situation, it looks like there is a collective requirement to save the three people. Among the three bystanders, the fair share required is derived. It seems clear that everyone here should save one person. This is best done by each co-operating to operate the winch. The CPB holds that none of the agents is required to sacrifice more than they would if everyone cooperated in that way.

However, Cullity then supposes that one of the bystanders shirks this responsibility. As the CPB holds that no one is required to sacrifice more than they would under full compliance (but they must do as much good as possible with that sacrifice), we might imagine that this amount of sacrifice could only save two of the three people. As Cullity notes, this does not seem right.

“This does not mean that the two of us who are left have done enough if we winch out only two people. Each of us should be prepared to save all three, given that the effort involved is small. And if we can only do so by cooperating, we should cooperate. The two of us now constitute a new group upon which a collective requirement falls – to help the people who still need it – and I should fully contribute to fulfilling this further requirement.”

(2004, p.77)

This looks even more problematic if we have introduced a case with elastic thresholds – where small additions can result in considerably more good. Many of the most pressing cases do not seem to be of this sort,<sup>21</sup> but we can easily imagine such cases.

Consider the following case:<sup>22</sup> A powerful and evil (but remarkably trustworthy) wizard arrives in a village, Agencia (a principality with a population of one hundred), claiming he will kill everyone in the neighbouring large city,

---

<sup>21</sup>In the case of world poverty it does seem to be the case that (within certain limits) the more money is given, the more lives can be saved, and that the amount of lives saved by one donation would be similar to the amount saved by an additional donation. There will be various thresholds, however, the passing of which will determine whether certain initiatives can take place.

<sup>22</sup>This is an elaborate fantasy case to emphasise the point. None are likely to be nearly as extreme in the actual world, but as discussed in section 1.1.1, we typically think moral theories should be universal, in that they should be able to give appropriate verdicts in all possible cases. Granted, wizards might not be possible, but this can be replaced with a powerful scientist if too objectionable!

Patiencia (population fifty thousand), unless he receives one thousand gold coins. The wizard tells them he will summon the one hundred villagers in alphabetical order into his tent to give him the gold, and if he does not have enough gold after the last person has left, he casts his evil spell. The villagers are all equally wealthy and can all afford to pay their fair share — ten gold coins — so at an impromptu village meeting decide to do so. Ada goes first, then the next, and so on. When the final villager, Zorgon, arrives in the tent, he pulls out his coin sack ready to hand over ten coins, when the wizard announces that he has only received 989 coins so far. Perhaps someone in Agencia has deliberately flouted the agreement, or maybe they have miscounted.

According to the CPB, Zorgon is not required to sacrifice any more than the ten coins. However, considering what is at stake, and the minimal additional cost, it would seem absurd to claim that Zorgon is not under an obligation to contribute eleven coins. Zorgon's doing the most good with *up to ten coins* might involve not giving the wizard any, but instead opting to buy a friend who is feeling a bit sad an ice cream later. The CPB might thus conclude that Zorgon must *either* buy a friend an ice cream *or* go "above and beyond" his duty and give the wizard eleven coins. The suggestion that buying the friend an ice cream would be permissible in this case seems ludicrous.

Though Murphy's view does take into account compliance, surely some other factors must be in play in determining our obligations. Though it isn't *fair* that some people might have to pick up the moral slack, this might still, in cases where the benefit of additional negligible sacrifice is enormous, be required of them. This would be evidenced by the blame felt (and appropriately felt) by the rest of the village towards Zorgon if he refused to pay the eleven coins to prevent the catastrophe.

Murphy acknowledges and accepts this conclusion (1993, p.290, n.43). Though it seems implausible, he notes that it looks like any principle except the Optimising Principle will be subject to structurally similar problems. At any limit, if it turns out that minimal additional sacrifice would result in some huge benefit, it will seem bizarre to claim that the additional sacrifice is not required. Murphy's response in the face of this problem is to admit that the CPB seems problematic, but suggest that it is still the best candidate we have.

### **Extreme Demands in Full Compliance**

Leaving aside the unusual verdicts that the CPB leads to in cases of partial compliance, we might object that it fails as a response to demandingness objections because not all instances of excessive demands seem to be caused by partial compliance. Murphy rejects this, and in doing so considers some extreme demands caused in other ways. One instance is cases of heroic self-sacrifice like the soldier who jumps on a grenade to save his comrades, and Captain Oates, the member of Scott's Antarctic expedition who left the tent to certain death when he realised his ill-health was jeopardising the chances that the others would survive (2000, p.99). Murphy notes that these acts, while requiring incredible bravery, do not actually entail sacrifice on his view, as in both cases the agent would have died anyway. Full compliance does require that they act in this way, but Murphy does not see this as a problem. We regard these acts as heroic because performing them is "motivationally very hard", but Murphy thinks the CPB is correct; that these agents *are* required to act as they do (2000, p.100).

As those cases do not involve what the CPB can regard as a sacrifice (because the agents have the same well-being whether they comply or not), they do not provide the best challenge. Murphy also discusses cases where one's well-being is affected:

"[S]uppose you are the only person in a position to prevent a nuclear accident that would kill many thousands of people. If you do what is needed, however, you will receive a painless but eventually fatal dose of radiation. Do you escape in the helicopter, leaving many thousands to die. or do you, in effect, sacrifice your life?"

(2000, p.99)

The CPB clearly requires the sacrifice here, and again Murphy accepts this verdict. This does seem very plausible, but by setting up the case in this way, with the benefit of complying so enormous, Murphy seems to have given an easy test. On any plausible weighting, Scheffler's agent centered prerogatives or Miller's *Sympathy* principle would also arrive at this verdict. In fact, any (even remotely plausible) account which accepts that there are duties of beneficence and that they can *sometimes* require one to sacrifice

one's life will come to this verdict. More tricky, would be if the preventing the incident would only save two people, or perhaps one person and a dog.<sup>23</sup> The CPB seems to still require the sacrifice, but this is *far* less intuitively obvious.

There also seem to be very different types of cases that can lead to excessive demands, even when compliance is full. Instances where the burden falls only on one agent provide the easiest cases, as full compliance is guaranteed simply by the individual complying. Consider, for example, Judith Jarvis Thomson's violinist example (1971, p.49). If, like Thomson, we think it would ask too much of an agent to require them to be bed-ridden for nine months even if this is the only way to preserve the life of a famous violinist, the CPB provides no way of responding to this.<sup>24</sup> Optimal behaviour would clearly involve making the sacrifice, and as the victim's compliance would ensure full compliance, the CPB would require that sacrifice.

For a more commonplace example, we might consider Fiona Woollard's demandingness objections to extremely burdensome prescriptions for pregnant women and new mothers (2016b). If we think the advice – and purported *duties* – for women to avoid not only foods like “unpasteurised cheese, raw eggs, pâté, raw meat, liver, etc.” but also “new cars, new furniture, air fresheners, deodorant, sunscreen, non-stick frying pans, and food in plastic containers” is too demanding, we cannot explain this away by considering rates of compliance.

The overly demanding prescriptions in these instances is high even under full compliance – because the compliance (or non-compliance) of others does not create additional ‘moral slack’ to be picked up – so any appeal to the Compliance Condition is of no help.

Murphy acknowledges that it would be a theoretical benefit to the Compliance Condition if adopting it dissolved demandingness objections and sees

---

<sup>23</sup>Metz provides cases structurally similar to this and makes a very similar objection, though his objection is clouded by his assumption that Murphy's argument is reliant upon utilitarian elements (2001, p.616). These problems are more serious than this, as they persist even if – as is the case in “The Demands of Beneficence” (2000, p.280)– the CPB is given in an axiologically neutral way.

<sup>24</sup>This case might look slightly different because the purported requirements violate some right to bodily autonomy, and we could suppose that this has a special status that provides an important moral distinction. If, however, the length of time in the violinist case was, say, a mere ten minutes, it seems obvious that one would be required to make that sacrifice. So, it seems that it is the *amount* demanded that makes the requirement intolerable, rather than the *type* of burden.

this as one reason to endorse it. He does, however, recognise that it might not be able to do this on its own. With that thought in mind, he suggests that it could be supplemented with a limiting principle making a *Limited Collective Principle of Beneficence* (1993, p.289). This would provide a constraint upon demands that meant that they could never be more burdensome than some extreme level. However, we can create cases like those provided above – where a principle makes a demand of only one agent – where the principle seems to demand too much, even though it demands an amount that *could be* required in certain cases. For example, consider Thomson’s violinist case once again. Even if we accept that in that case it would be good to volunteer but not morally required, our verdict is likely to change if the amount of people that could be saved was significantly higher. Being hooked up to a thousand people might seem implausible, but we could alter the example so it only(!) requires that one undergo extensive medical tests that restrict one’s movement to the same degree so that some rare antibodies could be collected from one’s blood. Presumably, in that instance, we would judge that the sacrifice is required, so a *Limited Collective Principle* would also fail.

Of course, it would be possible for Murphy to reply that even if his account fails to resolve overdemandingness in certain non-group cases, it might still succeed with the types of collective project he discusses.<sup>25</sup> But as overdemandingness was the starting position for Murphy’s discussion, and as Murphy saw the prospect of dissolving demandingness objections as a big motivation for suggesting the Compliance Condition, this would be a big weakness.

We might also think it desirable (if possible) to give a uniform explanation for why principles are overly demanding. Perhaps this could even extend to other normative domains, which share certain structural features, like rationality (e.g. Hedberg 2014) or etiquette (e.g. McElwee 2016a).<sup>26</sup> Given the non-group cases where demandingness seems to be a problem, it does not look like the CPB is able to give a uniform explanation even in the moral domain, and because of the agent-relative normative requirements in play, has no promise of responding in the other domains.

---

<sup>25</sup>Given the objection from small marginal sacrifices already discussed, I would be suspicious even of that claim.

<sup>26</sup>The notion of supererogation across normative domains is discussed further in section 5.2.4.

### 4.3 Explaining the failures

It seems that the recent attempts to respond to overdemandingness have serious flaws.<sup>27</sup> Not only do they give extremely counterintuitive verdicts in certain situations, but each has reasons why it seems unable to adequately respond to the *demandingness* itself. Miller’s attempt fails in this because it ties demands to *how well an agent’s life* is going (or going to go), not taking into account that demands might be too extensive, while keeping consumption of luxury goods the same. Murphy’s principle fails to account for demandingness specifically because it sees the problem of moral demandingness as a symptom of a failing to distribute demands when groups only partially comply. Although the paradigmatic cases – being required under a Singer-esque view of morality to donate practically all of one’s income to effective charities – are overly demanding because of partial compliance, this is not always a factor in generating the problem.

To better address the problem, I next suggest thinking about what is so bad about theories that make extremely burdensome prescriptions, particularly focussing on the unfair attributions of blameworthiness made by such theories. Thinking about when blame is and is not appropriate, and the relationship between blameworthiness and moral wrongness, may shine some light on what conditions must obtain for an agent to be subject to moral obligations.

---

<sup>27</sup>One notable recent attempt I have not discussed is Garrett Cullity’s. His solution in *The Moral Demands of Affluence* (2004) involves examining the presuppositions of beneficence. These presuppositions – facts that must be the case in order for the motivation to hold – provide limitations on beneficence. Particularly with regards to beneficence, Cullity claims it is a presupposition that an “interest in having or doing something can only ground a moral requirement. . . as long as it is not morally wrong for you to have or do that thing” (2009, pp.23-24). He then suggests that if it is permissible for us to help others pursue certain life goods, it must also be permissible for us to pursue them ourselves. While I do not have the space to fully discuss this here, I think Cullity’s account fails because it overgeneralises, in part because it does not adequately distinguish between needs and goals. Regrettably, I do not have space to discuss Cullity’s account adequately in this thesis.

## 5

# Blameworthiness, Difficulty and Obligation

*“Only in our virtues are we original, because virtue is difficult.”*

– Iris Murdoch, *Nuns and Soldiers*

*“ ‘What ho!’ I said. ‘What ho!’ said Motty. ‘What ho! What ho!’ ‘What ho! What ho! What ho!’ After that it seemed rather difficult to go on with the conversation.”*

– P.G. Wodehouse, *My Man Jeeves*

In chapter 4, I argued that there are serious problems of the weakened theory offered by Miller. Miller attempted, through his principle, to give verdicts of weaker obligations, or to permit sub-optimal actions in favour of one’s own interests. While he recognises that there is something problematic about extremely burdensome theories, he fails to say exactly what is the problem. Of course, as we discovered in chapter 3, this is no easy task. None of the obvious candidates (e.g. impartiality, maximisation) lend themselves as culprits for making a theory overly demanding.

Similarly, recent attempts that do make an effort to locate plausible problems of a theory that make them overly demanding, like Murphy’s,<sup>1</sup> also fail. Murphy (like Kwame Anthony Appiah ) thinks the problem consists in compliance (or lack thereof). As was suggested in the third chapter,

---

<sup>1</sup>I also regard Scheffler (1982) and Cullity’s (2004) comprehensive attempts in this regard.

compliance does seem to be a relevant feature, but Murphy's particular account seems problematic for the reasons given by Cullity.

In this chapter, we take a step back and consider obligation more abstractly. We begin by thinking about the notion of moral obligation and the closely related notion of blameworthiness. After this, I offer a solution to the demandingness objection based on recent work by Brian McElwee, considering the possibility that understanding obligations in terms of difficulty might be a way to more successfully account for the extent of the demands of morality. I reject this, but in doing so suggest that there may be a problem in the dialectic, which contains a presumption of doing the best. This leads us to problems when we consider theories susceptible to demandingness objections, because our everyday moral verdicts and practices contain no such presumption. I suggest that removing this shift and instead taking seriously considerations of how obligations are *generated*, many of our problems will be alleviated. If this dialectic shift is taken seriously and we consider deontic and axiological matters independently, I claim that we are unlikely to accept a theory which is overly demanding (if we do, the problem will be because of our considerations of deontic matters).

Of course, simply suggesting that we should address our deontic concerns somewhat differently is easy. The difficulty really arises in contemplating how we should think about our obligations if they are not – as some suppose – tightly linked to considerations of value. I examine intuitions in cases like Le Chambon, where we can see that obligations appear to be very context sensitive. I then discuss our moral phenomenology when we judge that we are under an obligation, and when we flout such an obligation.

Following this, I discuss what sort of success criteria are needed for a satisfactory theory of moral obligation (and what relation will be possible between obligation and value). These are proposed criteria by which we can judge whether a theory is successful. It should, for example, preserve our ability to (and make sense of our practices of) to partake in moral discourse. Ideally, we will also preserve notions of moral progress (e.g. that things are better now, without rampant slavery, than they were when the slave trade was in full flow). We presumably will also want to preserve the 'moderate' position that we do have moral obligations, and perhaps more substantial ones than many people regularly meet, but these are not as excessive as is claimed by the extremists.

## 5.1 Blameworthiness and Moral Wrongness

Clearly, demandingness objections concern moral obligation. They make the claim, of a certain moral theory, that its obligations are more burdensome than they should be. Definitionally, an agent who fails to meet an obligation (or commits an action that they were obligated not to commit) has done something morally wrong.

Some have claimed that moral wrongness is an elusive or intractable notion and might be better understood via some other concept.<sup>2</sup> We might think an easier concept to discuss is blameworthiness, and think that because of the close relationship between moral wrongness and blameworthiness, this can then help us in making some claims about obligation.

### 5.1.1 Some Thoughts About Blame

Before discussing blameworthiness more seriously, let us forestall one potential confusion. We should of course note that an action's being blameworthy (or an agent's being blameworthy for performing a certain action) does not mean that the *expression* of blaming is appropriate (or even possible) in the case.<sup>3</sup> Perhaps we can consider a case where a group is being driven to a destination. The driver claims that they have knowledge that one route is best, against the advice of several of the passengers. Because of his failure to take good advice, we would normally accept that the driver is blameworthy for the group's being late. This might appropriately be borne out in an expression of blame in some way. A passenger might rightly say "It's your fault we're so late" or "You should listen to us next time", or impose some other social censuring. However, perhaps the passengers know that the driver has recently had some sort of serious life problem, is currently finding it very difficult to deal with negative responses or is prone to extreme over-reactions that make him unable to perform simple tasks. In this instance, it might be inappropriate for them to blame him – that is, to express the blame sentiment. Obviously, in the latter case, it certainly would not be

---

<sup>2</sup>Derek Parfit, for example, suggests that something like this might be successful (2011, p.165). Richard Yetter Chappell agrees with this diagnosis and justifies his project in a similar way (2017).

Mulgan actually suggests the contrary claim — that our strongest relate to "the rightness and wrongness of actions" (2006, p.3).

<sup>3</sup>Skorupski takes pains to clarify this distinction in *The Domain of Reasons*, and suggests that this is an area Mill was insufficiently clear on (2011, p.292).

appropriate to blame him while he driving. This distinction – between the blame sentiment and the act of expressing blame – is important to keep in mind. When someone is *blameworthy* this can be understood as their being *blame-feeling-worthy* (Skorupski, 2011, p.292).

So, what is it to *blame* someone? While we may say in ordinary language that something, even inanimate objects, are ‘to blame’ for some state of affairs – like “the faulty brakes were to blame for the crash” or “I blame the bad weather!” – our talk of moral blame is more than simply establishing a causal role. It is also clear that blame can be rightly felt (or appropriate) and yet not expressed. This much is generally agreed upon, but the exact nature of the role of blame in our moral lives is not.

One feature of blame is that an instance of blame makes a claim of a moral failing. For it to be appropriate to feel blame towards someone for being late, doing so must involve a negative evaluation of her conduct. However, Smart correctly noted that our notion of blame is more than that of moral grading (1961, p.303). While there is a way of thinking about praise as simple evaluation, the purpose of which is communicative — to indicate that someone possesses a skill or has performed in such a way — blame doesn’t work this way. Saying that someone is a very good singer is an act of praising, while saying that they are not good at singing is (generally) not one of blaming.

It is a contested issue what the *proper* function of blame is. Macalester Bell identifies five distinct aims of blame, each of which are valuable in a certain way.

1. Blame marks the damage done to our relationships.
2. Blame informs wrongdoers of their wrongdoing.
3. Blame motivates its target.
4. Blame educates and motivates the moral community.
5. Blame expresses condemnation / avoids moral condonation.

(2013, pp.267-268)

These are all plausible effects that acts of blaming might have. Clearly not all expressions of blame will involve *all* of these features, or even have

all these features as possible outcomes. It may not be possible to damage a relationship that is non-existent or already toxic.<sup>4</sup> We need not accept that these effects must all (or even mostly) be present for an expression to count as blaming, but these are often consequences of acts of blame.

Scanlon sees the first these as the most important. His account of blame focusses on how one's relationship with the target of the blame is altered. He claims that "to blame a person is to judge him or her to be blameworthy and to take your relationship with him or her to be modified in a way that this judgment of impaired relations holds to be appropriate" (2008, pp.128-129). It also requires that the target is a moral agent, the type of relationship one must be engaged in for blame to be appropriate is a "moral relationship that we stand in to all rational creatures" (2013, p.87). Sher, providing a different type of account, sees blaming as including "a set of affective and behavioral dispositions, each of which can be traced to a single desire that the person had not performed his past act or not have his current bad character" (2006, p.112). Strawson famously regards blame as an expression of a reactive attitude. This expression (if successful) will typically have some of the effects Bell describes. He also notes that blaming requires treating a person as a moral agent. Strawson expresses this thought when he contrasts the attitude involved in a human relationship with "the objective attitude", in which someone is regarded as merely "an object of social policy", something "to be managed or handled or cured or trained" (2008, p.9).

While Strawson, Scanlon and Sher offer very different accounts of what it is to blame, they share certain features. Expressing blame is only appropriate in the context of moral agents. When blame is applied, it implies that the target has exhibited some moral failing. Because of its social roles, blame seems to be a concept extremely important to our navigating the moral world and our living with others generally.

Due to of the variety of negative consequences of being blamed for an individual, we do not want to be the appropriate targets of blame, i.e. we do not want to be blameworthy.

---

<sup>4</sup>Because Scanlon sees blame as involving modification of relationships, this type of case offers a prominent problem for his account of blame – the 'stranger problem' (e.g. 2013).

### 5.1.2 The Relationship Between Blameworthiness and Wrongness

So far in this chapter, I have mentioned the close connection between blameworthiness and moral wrongness, but given no suggestions as to the nature of this relationship. One natural thought that we might have is that the blame sentiment is only appropriate when a morally wrong action has been committed, or more formally:

If  $\phi$ -ing is blameworthy, then  $\phi$ -ing is morally wrong.

This seems to be very plausible.<sup>5</sup> With that in mind, should we accept that every blameworthy act is morally wrong? There are two types of responses that suggest not. The first involves what Julia Driver calls suberogation.

#### Suberogation

Driver introduced this term to describe a category of acts purported to be blameworthy but not wrong, or bad to do but not forbidden (1992). Driver provided three types of examples she supposed could fit in this category: ‘morally charged’ situations, the case of multiple abortions, and owing favours.

Morally charged situations are ones in which it is permissible to do two options, but one is good, while the other is bad. Driver gives examples of a man refusing to swap seats to accommodate an elderly couple on a train, and a man who can donate a kidney to save his brother (1992, pp.286-7). In each case, Driver suggests that while it would be praiseworthy to act well, failing to do so is blameworthy but not wrong. She suggests that as the agents are within their rights not to do the good action, it cannot be obligatory, so failing to do so cannot be wrong.

I suggest that Driver is mistaken about the correct verdicts in these cases. In some instances, she seems to be assuming that just because someone acts ‘within his rights’ that he does not act wrongly in doing so. In *The Merchant of Venice*, Shylock gains a legal right to a pound of Antonio’s flesh, but it would still be wrong to him to invoke that right. In the seating instance, it

---

<sup>5</sup>This also seems to be the standard position, e.g. Gibbard (1990, p.42), Skorupski (2011), McElwee (2016b, p.21), (Bell 2013, p.265), McGeer (2013, p.162).

seems plausible that it is just wrong to not move. In the kidney case, Driver's judgment that the refusal to donate is blameworthy also seems suspect. She suggests that "failing to act on an ideal does seem blameworthy" (1992, p.288), but this simply seems mistaken. If this was the standard below which we deemed people blameworthy, everyone would (almost constantly) fall into this category. The notion that everyone – including people who spend impressive amounts of time and money trying to make the world a better place – merits blame is precisely the type of thought which makes extremely demanding theories seem so implausible.

The case of multiple abortions is a peculiar one. Driver insists that the abortions must be 'frivolous' in order to be suberogatory (1992, p.293). It is unclear exactly what the frivolity is thought to indicate here. If they are truly accepted as frivolous (and there is any value associated with the life of the foetus), even one seems blameworthy, and suited to a similar diagnosis as the morally charged cases. Previously, however, she presented the problem as one with a different structure:

"If abortion is permissible, then multiple abortions are permissible. But a woman who has nine abortions has done something morally questionable...if every single abortion is unreservedly permissible, why this feeling that there's something wrong with the eighth, the ninth or tenth (per agent)?"

(1992, pp.288-289)

This problem seems to call to attention problematic character traits, but with many parallel cases. We might imagine the decision to give (or not give) money to charity. If an affluent and informed agent goes one month without donating any money towards relieving some of the vast suffering in the world, this might be permissible, but if they never donate any money, they are likely to be judged as having morally failed. Structurally similar examples might be extremely mundane. If you lived with several other people, and there was a chore that needed performing each week, not doing it on any individual week is permissible, but deliberately shirking regularly would not be.

The question of how many of this type of action makes an agent blameworthy is a coarse-grained question, but once the notion that one cannot

commit wrongs by acting within one's rights is dispensed of, this type of example offers no particular problems of its own. We can simply say that agents *do* act wrongly when they act in this way – even if we judge that we must allow them to do so.

The final example Driver makes use of involves favours. If one agent has helped another you on many occasions in the past, ordinary moral thinking would suggest that you 'owed them one'. Driver again seems convinced that this could not be obligatory – she says the agent failing to deliver on a favour owed "is not doing anything wrong" (1992, p.289). It is unclear why she thinks this. It might again be that no agent has a right to having favours returned, but just as it would be wrong to let a child drown in Peter Singer's pond, regardless of whether we think there is a right to be saved from ponds, one can act wrongly in this sort of case, even if there is no right to favours.

For the reasons provided above, I see no reason to accept Driver's claims that suberogatory acts do exist. I would posit instead, that the acts in question either need to be understood as wrong when taken as a whole (like not doing one's part to keep your residence clean) or, as perhaps wrong in each individual case. Even if each individual "frivolous abortion" or "not cleaning of the kitchen" is wrong, we might judge that they are the types of cases where it is not appropriate to express blame – even though they are blameworthy.

### **An Epistemic Dimension of Blame?**

A second way we might think that an act might be blameworthy but not morally wrong is if we accept that blameworthiness depends on the epistemic position of the agent, but that moral wrongness depends on effects or facts. We could then give a paradigmatic case of wrongdoing wherein an agent is also blameworthy, but then tweak the example such that either the agent's epistemic position is such that they could not be aware that any negative effects would occur, or so that no negative effects actually obtain. This position (and this method of elucidation), is advocated by Ishtiyaque Haji. He considers cases involving two doctors to elucidate his position. In these cases he attempts to give cases of blameless wrongdoing and blameworthiness without wrongness.

The first case – *Howser's Hell* – Doctor Howser, does everything she can to save her patient, and from the diagnosis she had come to (with good

reason) gives her patient medicine A. Unfortunately, the diagnosis, through no fault of Doctor Howser, was incorrect (1997, p.528). Haji suggests that this is a case of blameless wrongdoing .

The second case (*Deadly's Defeat*) involves Doctor Deadly, who (unsurprisingly) tries to kill her patient with some medicine that will accomplish that according to the diagnosis, but is foiled by bad luck. It just so happens that due to some rare and unnoticed fact about the patient's condition, this medicine is just what was needed, and the patient makes a full recovery. Haji suggests that although injecting the patient with the medicine here was the right thing to do, Deadly is to blameworthy for her conduct. Dr Deadly is clearly the problematic case for the above claim, as this purports to provide a case wherein an action is blameworthy but not wrong.

While I accept Haji's claim that *Deadly's Defeat* gives us a case where the agent is blameworthy, it is still not clear that it provides a cases without wrongdoing.<sup>6</sup> In the remainder of this section, I will sketch out some points in favour of rejecting that cases of blameworthiness without moral wrongness are possible.<sup>7</sup>

There are two types of manoeuvres one might use to defend the above conditional against cases like this. Firstly, one might hold that actions can only be morally evaluated under certain descriptions. If, for instance, only actions under descriptions that could be salient to agents at the time of performance were morally evaluable, we could say that Doctor Deadly acted rightly in trying to kill the patient, but that saving the patient was not something she could be morally evaluated for. What the agent really did, in this sense, is what Austin would refer to as what the agent *baldly* did. We may...

“argue that it is not quite fair or correct to say baldly ‘X did A’. We might say it isn't fair just to say X did it; perhaps he was under somebody's influence, or was nudged. Or, it isn't fair to say baldly he did A; it may have been partly accidental, or an intentional slip. Or it isn't fair to say he did simply A — he was

---

<sup>6</sup>Parfit discusses cases like this, but distinguishes between what is “objectively right” and “subjectively right”, the latter of which depends on one's epistemic position (1984, p.24).

<sup>7</sup>It might be noted that the evaluation of rightness/wrongness in play here is a very consequentialist one; what has the best consequences is regarded as the right action. While it could be argued that this is a mistake, I will not press that claim here.

really doing something quite different and A was only incidental, or he was looking at the whole thing quite differently.”

(1961, p.124)

Perhaps we would be tempted to impute moral responsibility only for actions it might be properly said that an agent ‘baldly’ did. This type of manoeuvre would require major development and revision, and while I think it may offer promising results, defending such a radical manoeuvre is not required here.

A second approach is motivated by the thought that Haji’s proposal struggles to make sense of claims that it is wrong to attempt to perform wrongs. When Doctor Deadly tries to kill her patient, Haji claims she is only blameworthy. But what is she blameworthy *for*?

It would serve purposes of discussion here to discuss why Haji thinks that blame should be so divorced from wrongness. Haji motivates his position via bringing out a tension that arises between two commonly-held principles:

**Objective Wrongness** S has a moral obligation to perform [not to perform] A if and only if it is morally wrong for S not to perform [to perform] A;

and

**Ought Implies Can** S has a moral obligation to perform [not to perform] A only if it is within S’s power to perform [not to perform] A.

(Haji, 1997, p.526)

If we accept these principles, Haji claims, we are committed to the claim that it “is morally wrong for S to perform [not to perform] A only if it is within S’s power not to perform [to perform] A”, the claim that wrongness depends on alternative possibilities, which Haji labels (WAP) (1997, p.527). Haji sees this as plausible, but as forcing us to deny that blameworthiness depends on wrongness, because of verdicts it provides in Frankfurt-style examples.

In a typical Frankfurt-style case, an agent harbours an intention to commit some wrong,  $\phi$  at time  $t$ . Unbeknownst to them, if they display any indication of not  $\phi$ -ing at  $t$ , a counterfactual intervener (often a wizard, or

ingenious neuroscientist) will make it the case that the agent does actually  $\phi$  (1998a).

Because of the Frankfurt-style cases, while we might accept that *wrongness* depends on alternative possibilities, we certainly do not want to accept that blameworthiness does. If we did accept that, then the agent who *does* freely choose to assassinate someone, with the counterfactual intervener lurking in the background, would be blameless. Frankfurt's cases aim to preserve moral responsibility in these cases.

Haji claims that in these cases, the agent does nothing *wrong*, but they still do something blameworthy. But what is it that these agents do that is blameworthy? In his analysis, Haji claims that when Doctor Deadly injects the patient with the cure, she is blameworthy for *injecting the patient* — it is the physical behaviour for which the agent is blameworthy.

This leads to some strange results. On Haji's account, Doctor Deadly is blameworthy for injecting the patient (even though it saved her), while Doctor Howser is praiseworthy for injecting the patient (even though it killed her).

A much more natural way to present the blame here would be that Doctor Deadly was blameworthy for *attempting* to kill the patient. Similarly, Doctor Howser is praiseworthy for *intending* to save the patient.

Once we have reached this evaluation, however, we can ask whether what those agents were blameworthy for was also morally wrong. Either, one could suggest that these are mental actions, and thereby preserve the notion that the concepts of wrongness and blameworthiness both predicate over *actions*. Alternatively, one could make the case that something like intentions might not count as full-blown actions, but might still be the type of things that can be wrong/blameworthy.

Either of these manoeuvres has some plausibility. While Frankfurt suggests that actions require a purposive guided *movement* (1998b), others consider mental actions to be an obvious and ordinary type of action. Mele, for example holds that *deciding to act* — what he calls practical deciding (distinguished from cognitive deciding, which is simply deciding *that* something is the case) — is a type of action (2003, p.198).

If, on the other hand, we decide that intentions or decisions are not actions, this is still not a death-knell for the view that something (whatever the relevant category is) must be morally wrong in order to be blameworthy.

Lifting the requirement that wrongness must be a property attributed only to the domain of actions might actually make some cases easier to explain. If we consider cases of culpable ignorance, for example, it can be very difficult to locate a specific action that was at fault. Forgetting may be another such example.

I do not take a stance on which of these routes is likely to be best here, though as in 1.1.3, I will continue to refer to decisions, tryings, volitions and the like as actions. This is intended only as a loose way of talking, rather than a serious commitment to the existence of mental actions. What I take to be important here is that it can be wrong (as well as blameworthy) of an agent to make a certain *decision* (regardless of its action-status).

While I have only briefly sketched some of these arguments, there are several ways one can defend against Haji's claim, and these give results more in keeping with our intuitions about what people are blameworthy for. More can certainly be said on these matters, but more substantial discussion is not required here. For our purposes here, I will simply accept the conditional "If agent A is blameworthy for  $\phi$ -ing, then it must be true that it was morally wrong of A to  $\phi$ ".

### 5.1.3 Blame and Demandingness

If what has been argued for above holds, this makes the set of blameworthy acts a subset of the set of morally wrong acts. It should be noted that the set of blameworthy acts is also a proper subset of the set of morally wrong acts, as there are wrong acts which do not seem to be blameworthy. Haji's *Howser's Hell* case purports to give one such example, where the doctor gives her patient what she takes to be drug that will cure, but, unbeknownst to her and through no fault of her own, the drug actually kills the patient.<sup>8</sup>

As well as cases where the agent's epistemic situation seems to alleviate blame, there are also cases where agents are not judged as blameworthy even when they do know all the details of a situation. Let us consider a man who has told his wife that he will cook dinner that evening. This is

---

<sup>8</sup>If one of the above strategies is taken, this actually will not provide an adequate example. If one holds that only descriptions which could be salient to the agent as something she *does* at the time of the behaviour can be regarded as actions, and that only actions can be judged as wrong, *Howser's Hell* will not provide such a case. We might see this as in keeping with Austin's analysis of what an agent 'baldly' does. However, the following examples do not share that feature, so I need not dwell on this point.

a fairly obvious case where this speech act confers upon him an obligation to perform an act – namely cook dinner that evening. We might imagine situations obtaining however that make the day much more stressful and exhausting for the man than he could have anticipated. Perhaps his mother calls with a host of problems, the dog gets ill and needs to be taken to the vet. We might think that the obligation is defeated if the events are too exhausting (I’m staying neutral on this), or if they take up enough time to make it impossible to also make dinner, but it looks like there is an amount of stress/exhaustion that would make it the case that he ought to make dinner, but that – given the events – it would be excusable not to do so. If, when his wife arrives home, he tells her of the stressful day (and apologises for not cooking), it would be inappropriate for her to blame him (or even *feel blame* towards him) for not making dinner.

This discussion of blameworthiness was motivated by the thought that blame is a more tractable notion than moral wrongness. It seems clear that we want our moral theory to forbid actions that are blameworthy. We should also be aware that blame is a primary motivation for raising demandingness objections. We find it extremely inappropriate to say of some agents that their conduct is morally wrong. Consider McElwee’s example of John, a typical first-world, reasonably well-paid agent, who, after being struck by the extent of suffering in the world. . .

“... resolves to donate 30% of his income to Charity X. He also spends two nights a week raising money for Charity X, writing letters to newspapers and politicians campaigning for greater justice in international trade, and organising campaigns to raise awareness of unfair and harmful trade conditions. But John does less than the most he can to produce the best available consequences. He could be giving up, say, 75% of his income, and six nights a week to good causes (while still retaining his psychological health and the ability to carry on this way of life over a long period of time).”

(2010, p.395)

If we accept an extremely demanding moral theory – one which claims he morally ought to be giving 75% and donating six nights a week – this gives

the verdict that John's conduct was *wrong*. One of the reasons this sounds so implausible is that there is an implication of condemnation, and further, an implication of blame that we find troubling here. Now, we might suggest, as many maximising consequentialists would be likely to, that though he does act wrongly, he is not blameworthy because in this sort of case he has an excuse.

Worth serious consideration are the cases where we do not deem blame appropriate, and whether in those sorts of cases it *is* a matter of excuses that make this the case, or rather whether something about the situation determines that the action actually isn't wrong. Those making demandingness objections seemingly have to accept the latter claim. One response that has been raised recently is that the difficulty of an action seems able to alleviate the blameworthiness of an agent who fails to perform it, and that because of this we should consider whether the difficulty of an action can make it non-obligatory. This discussion will be the focus of the following section.

## 5.2 Difficulty

Until recently, difficulty has been largely neglected as a potential explanation for a moral theory's overdemandingness. This might strike us as surprising, considering how definitions of 'difficult' are very similar to definitions of 'demanding'.<sup>9</sup> Recently, a few authors have discussed whether or how the difficulty of an action may affect its moral status.<sup>10</sup> Just how difficulty relates to moral obligation is not at all obvious. Amongst the few philosophers discussing the subject there is already significant disagreement. One natural thought we may have, which was mooted towards the end of the previous section, is that if an action is particularly difficult for an agent, it is less likely to be morally obligatory. Even more boldly, we might suppose that demandingness objections to moral theories *really* point to the difficulty of the actions they prescribe as problematic.

One motivation for viewing difficulty as having this role is that it could explain why certain factors are of normative importance. For instance, when making his case for the moral equivalence of watching a child drown and

---

<sup>9</sup>This point has been noted by McElwee (2017, p.23).

<sup>10</sup>See McElwee (2016b), Gustafsson (2016), Cohen (2000a) and Chappell (2017).

failing to donate some amount that would save a child to a charity, Singer claims that both the number of people not acting and the proximity to the child are morally irrelevant (1972, p.232, 233). If an action's obligatoriness is accounted for, even in part, by difficulty, we can see why these factors can be relevant. If no one else acts when confronted by a tragedy, it makes it easier to not act yourself. Contrarily, if everyone else is acting (jumping in to save children or donating), it might be more psychologically difficult to omit to act. Proximity, similarly, can make it more difficult to ignore someone's suffering. We could thereby explain the moral relevance of social or psychological factors.

This explanation would also allow us to make a concession to Singer. We could admit that proximity to the victims or numbers of onlookers are not *directly* morally relevant. This captures something important of our intuitions, as these factors do not seem to matter. Why should it matter how close someone is if you can help them just as easily? A difficulty-based account of demandingness could say that Singer is right that these features do not matter *directly*, or in themselves. Rather, these factors are one way for an action to be more difficult, and it is quite plausible (for independent reasons) that the *difficulty* of an action can affect whether or not it is obligatory.

In chapter 1 (and again when discussing Scheffler and Miller in chapter 4), we discussed the problems with seeing a theory as too demanding because it imposes too many costs. Some act's being more difficult, however, does not require that it entails any additional costs. McElwee suggests that it is problematic to conflate difficulty and costliness (and notes that it is often assumed that demandingness concerns only the latter (2016b)). The problems of this conflation are made salient when we consider what distinguishes the concepts. G. A. Cohen tries to elucidate the distinction between cost and difficulty through an example of a person who needs to help his friend get to airport:

“A friend desperately needs to get to the airport, and I can secure his presence thereon time either by giving him ten pounds for a taxi fare, or by driving him there on my bicycle, with him sitting on the bar. I am poor, but I have ten pounds in my pocket at my fingertips, so it is not difficult but it is costly for me to give him the money for the taxi fare. It is costly because

giving ten pounds involves a large sacrifice for me, but it is not difficult for me to supply him with it, as it would be if I had to go to a bank the other side of London to get the money.”

(2000b, pp.238-239)

While Cohen clearly picks out something important with this distinction, difficulty is also going to be a factor in the case of giving the money. We can easily imagine someone, after realising that giving £10 is the only way to help their friend, really struggling to give them the money. This of course isn't a physical difficulty — putting your hand in your pocket to reach for a £10 note is, for most of us, a relatively simple matter — but it is a difficulty nonetheless. We might see it as a *motivational* difficulty.<sup>11</sup>

Elsewhere, Cohen even begrudgingly acknowledges this type of difficulty (2000a, p.24), but claims that this sort of difficulty cannot serve to stop something being morally required. He claims that though it is “unreasonable to ask someone to do something impossible... it's not unreasonable to ask them to do something very difficult”.

Cohen does little to justify this claim. He discusses the distinction in relation to the difference between voluntarily donating some money to the world's poor, and having that same money taken by a government and directed for the same purpose. When thinking about writing the cheque, Cohen suggests that the difficulty in “contemplating the consequent cost can generate a cost not in the throes of the decision itself”, but doesn't see this as significant, and even then sees the difficulty as only contributing insofar as it can increase the costs (2000a, p.25).

One issue arises in the variety of ways we can understand the notion of ‘difficulty’, or the different ways something might be ‘difficult’. Consider, for example, hitting a bullseye on a dartboard. It seems natural to say that this would be (for most of us) difficult. ‘Difficult’ here seems to relate to the likelihood of success. However, it does not seem to mean only likelihood of success. It would be infelicitous to say of some statistically improbable events that they would be difficult. For example, winning the lottery is extremely unlikely. Similarly, the chance for any individual agent, that the sperm cell would win out and successfully fertilise the egg to make an embryo

---

<sup>11</sup>McElwee uses this term (2016b, p.25). Chappell, in referring to a similar phenomenon uses the terms *mental effort* or exertion of *willpower* (2017).

that became them, is extremely slim. However, despite the low probabilities of an agent succeeding in these tasks, it would not seem appropriate to call them difficult. What seems relevant in the distinction here is that deliberately hitting a bullseye requires that the agent orients themselves towards that goal, whereas winning the lottery (in an important sense) or being born are not things one can try to do.

In other situations the effort required (and duration) is what makes something difficult. Being a parent is difficult, for example. Here it clearly isn't about the probability of success — most people in the history of the world have done this (and not simply in the sense of successfully reproducing) — but the long-term commitment that is required.

Motivational difficulty can relate to both of these. If you are unlikely to succeed in an activity, it can be very difficult to motivate yourself to continue. Applying for jobs in a horrific job market, for example, might make an agent wonder, “What’s the point?” and make the process difficult. It also seems like it can be difficult to get oneself to commit to something which you could guarantee success in, like a long-term project or an increased monthly donation. This is the difficulty in getting oneself to act, which is importantly separate from the costs themselves.

It is often (if not always) impossible to completely separate costs from difficulty. This is because the things that can make something difficult — spending effort, time, money — also impose a cost. McElwee notes that we can consider cases wherein the best thing to do *for the agent* comes from a particular difficult action, in order to attempt to pull apart difficulty and costs (2016b, p.26). For example, we might consider an agent thinking about undertaking a strict exercise regime. The agent might realise that doing so would be the best thing they could do — it might make them healthier, more attractive and generally happier — so despite additional difficulty, would result in a huge welfare gain. These cases could be useful because, in a sense, there are no costs for the agent — they would actually benefit from the action. In this type of case, however, we can see that there are welfare costs, even if they would not be ‘net costs’. For example, we might consider the burdens as costs to a person at a certain time slice.

Because difficulty is usually (if not always) inextricable from some cost, it makes analysing its moral relevance problematic, especially if we want to contrast its moral relevance with that of costs.

One type of example that we might consider is of an agent suffering from acute depression.<sup>12</sup> Some days, he finds mundane activities extremely easy. On other days, summoning the effort to even get out of bed seems impossible, and he often fails because of this. While in his bed, we may imagine, that the agent realises the costs/benefits accurately both on the good day and the bad one – they seem to be the same. (This might be dubious, as being depressed is likely to deplete the enjoyment of ordinary tasks, but it seems like we can imagine an agent surmising that the costs/benefits are the same on these days.) The agent might acknowledge, in this situation, that the costs of remaining in bed are the same on each day, but on a bad day rousing the effort is just more difficult.

We can try also to highlight motivational difficulty in a case like Cohen's donation example. To focus on it, we can contrast with a case the same in all relevant respects. Consider Bill and Bob. Bill is filling out his tax return as he usually does, and discovers on the new form that his government has introduced a temporary additional tax to relieve some emergency, and is required to tick a box to illustrate his understanding of the situation. Bob is in the same situation, however his country's government has decided to present their citizens with the option of paying this money as a donation, and to choose to do so, citizens must simply tick the box.

Bill reads the description, sighs, and ticks the box. The cost has simply been exacted from him. Bob, however, reads the option, pauses, and perhaps considers the small kitchen renovation he will have to put off, or luxury good he will not be able to afford any time soon, if he does tick the box. The internal strife that takes place – or may take place if Bill genuinely thinks that the donation would do more good to others than if he kept the money for himself – is the motivational difficulty. Of course, going through this process may be stressful or unpleasant, so may also impose some costs to his wellbeing. When considering whether a potential action's high difficulty may in itself make it permissible to omit, those costs (the ones to the agent's welfare) must be ignored.

Could the sheer difficulty that an agent would undergo if they performed the action considered make it okay for them not to do it? Cohen rejects

---

<sup>12</sup>Chappell gives a structurally similar example of a lazy agent, who knows “full well that they will feel better - even in the moment - if they do it” and accepts that the “only ‘unpleasantness’ here is in the prospect of bringing themselves to actually do it” (2017, p.8).

this, claiming that an action's difficulty "is no reason for not performing an action that (although difficult) is possible" (2000a, p.24). He deems difficulty as irrelevant to moral obligation, but there are good reasons to think it is important. In the next section, I will provide an argument for giving difficulty a larger role.

We might consider two claims, a strong thesis and a weak thesis:

**Strong Difficulty Thesis** The demandingness of a moral theory is fully accounted for by the difficulty of its prescriptions; if a theory is objectionably demanding, this is explained by its prescriptions being too difficult.

**Weak Difficulty Thesis** Demandingness is partly accounted for in terms of the difficulty of the prescriptions; if a theory is objectionably demanding, this may be due to the prescriptions being too difficult.

It should also be clarified that both these theses concern difficulty *directly*. If an agent is strongly considering two alternatives, one of which benefits them, but the other of which there is more moral reason to do, the associated difficulty may result in some mental anguish, which could be understood as a cost to the agent's welfare. Any modification of what is obligatory – if this is deemed relevant at all – from this anguish can be seen as an *indirect* result of the difficulty. To elucidate, consider hedonic utilitarianism, suggesting that *all that matters* is pleasure. While the difficulty may reduce the agent's pleasure here, and plausibly result in changes in the hedonic calculus, it only matters *derivatively*.

The argument presented 5.2.1 aims to establish the weak thesis, though with a little supplementary argument could be used to endorse the strong thesis.

### 5.2.1 An Argument from Blame

This subsection employs the notion of blameworthiness to form an argument for the Weak Difficulty Thesis. One motivation for proceeding in this way is the thought mentioned earlier that moral wrongness is an elusive or intractable notion, so it may be better understood via some other concept. I will present two generally accepted claims, and put forward the case that the best explanation for these is the Weak Difficulty Thesis.

Consider two phenomena we may consider worthy of explanation:

1. We suppose that being *too demanding* is a reason to reject a moral theory.

That this claim is true is evident by the mass of literature regarding the demandingness of morality. That we are correct to suppose this is essentially what I argue for in chapter 2, in rejecting the extremist positions that dismiss the objection. If we think that some moral theories should be rejected simply because of how demanding they are — or because they could ask too much<sup>13</sup> — of some moral agents, we accept this claim. I shall take this as given.

For the reasons discussed in 2.1, appealing simply to costs imposed cannot explain a theory's being too demanding. Now consider the second claim:

2. Our blaming (and praising) practices *do*, and appropriately, make allowances for difficulty.

I take at very least the descriptive component of this claim to be fairly uncontroversial. To be clear, I intend this claim to include shifts in blaming/praising in both directions, i.e. making an otherwise praiseworthy act less or non-praiseworthy (e.g. if it is extremely easy in a particular circumstance) or making an otherwise blameworthy act less or non-blameworthy (e.g. if it would be extremely difficult to refrain from doing it in this particular circumstance).

In practice, if an agent fails to perform a desirable action in a certain situation, but further contextual features come to light which make it clear that performing that action would have been particularly difficult for that agent, we are likely to find them less (or not at all) blameworthy.

Consider, for example, the agent suffering from depression mentioned earlier, and a colleague casting verdicts about them. Though the colleague might originally deem him blameworthy for, say, missing an unimportant meeting, the discovery that his depression made it extremely difficult for him to motivate himself to perform even ordinary tasks, might cause them

---

<sup>13</sup>I add this caveat because we may imagine a moral theory which we deem appropriately demanding given the current state of affairs, but which may seem too demanding if, for example, much/all of the easily relievable suffering has been eliminated. Ashford makes the observation that Scanlon-style contractualism is extremely demanding in any practically realisable state of the world (2003). This may still serve as cause to reject a given theory.

to rescind that judgment, or at least make them judge the conduct less blameworthy.

More controversial is the question of whether it is appropriate that our verdicts of blameworthiness (and praiseworthiness) vary in this way. In considering this, it is useful to think back to what we have said (in 5.1) about the role or function of blame. One feature of blame is that the blaming agent will be disposed to modify her evaluation of the blamed person's character. Would we accept that we should differently modify (or not at all modify) our evaluation of a person for not  $\phi$ -ing, if we discover that,  $\phi$ -ing, for her, would have been extremely difficult?

It seems at least plausible that the difficulty experienced by the agent is relevant here. It makes sense that our praising and blaming practices, having a significant role in our actual social interactions, would be catered to account for human failings or weaknesses. Perhaps some accounts of blame and blameworthiness could make defending this claim more troublesome, but I regard this as fairly uncontentious.

Accepting these two claims as true, we might hypothesise that the best explanation for both these claims is that the permissibility (or impermissibility) of an action can depend (at least in part) upon its difficulty (i.e. that the Weak Difficulty Thesis is true). If then, we accept the truth of these claims, and that the best explanation for them is the Weak Difficulty Thesis, we should also accept this thesis.

The conclusion to this argument aims at the Weak Difficulty Thesis. This is tentatively defended by McElwee, though he rejects the Strong Difficulty Thesis. After arguing that “considerations of difficulty appear independently relevant in determining the extent of appropriate moral demands” (2016b, p.31), he argues that *costs to the agent* are also important, even without reference to difficulty. In attempting to support this position, he makes the following claim:

It seems plausible to say that the simple fact that performing some act, A1, will lead to a cost to me, while some alternative act, A2, will lead to a larger cost falling instead upon someone else, can be sufficient to ground a moral permission to perform A2 instead of A1, *even if somehow performing A1 is as easy as performing A2.*

(2016b, p.31)

A few thoughts about this type of example should be mentioned. McElwee notes that it is “difficult to have firm intuitions” about it because “given the motivational structures humans almost universally have, imposing a cost upon oneself is almost invariably more difficult than imposing a cost on a stranger” (2016b, p.31). This is certainly a difficulty we might have here. Another consideration that we should note is that *even if* McElwee is correct for some assignment of costs in A1 and A2 (which again is hard to say because of the nature of the case), it definitely isn’t the case for all costs meeting these descriptions. If you could accept a very small cost on yourself in A1 or impose a terrible one on someone else in A2, this definitely does not seem to ground any permission to perform A2.

Contra McElwee (and Cohen), we might also speculate that costs can be disregarded entirely. Against the inclusion of costs in demandingness at all, I would offer several points. Firstly, it looks like costs will be accounted for in any plausible moral theory. Understanding the costs of an action as Cohen does, as “what I would lose (but would have preferred to keep) as the result of performing it” (2000a, p.23), these include losses of possessions, abilities, health, time, opportunities and all manner of uncontroversially good things. They are losses to interests, and any theory which cannot account for their goodness (or that they provide a pro tanto reason in favour of some course of action), and factor these into certain moral decisions, is a theory that has failed independently of its understanding of demandingness. If we are already involving difficulty in our understanding of demandingness, then this includes the difficulty of imposing costs on oneself, so we might wonder what work costs are doing.

That this is the case does nothing to undermine McElwee’s case above, as it stipulates that there is no difference in difficulty between the cases, but it is important to keep in mind that costs are being counted somehow, as costs to an agent (just like costs to a moral patient) do enter one’s moral evaluations. Why then, should we expect costs to feature as an aspect of demandingness at all?

To spell this out further, here is the claim: high costs to an agent cannot make an action too demanding, unless these costs outweigh the costs to others,<sup>14</sup> in which case *the action will not be morally required anyway*.

---

<sup>14</sup>This ‘outweighing’ may be understood in various ways. We might think that the

I have little more to say here, but it might be suggested – if we accept that costs are ruled out, and no other factors seem able to contribute to shifting obligations – that McElwee should be committed to the Strong Difficulty Thesis. On the other hand, there may be independent reasons for thinking that costs to an agent matter, or additional considerations of import.

Regardless, for the reasons given in the following subsection, I will ultimately reject both these positions.

### 5.2.2 Difficulty as Merely an Excuse Condition

It is generally accepted, as discussed earlier, that a morally wrong action is blameworthy except when an extenuating circumstance applies.<sup>15</sup> Extenuating circumstances include justifications and excuses.<sup>16</sup> To say that an action is justified “is to say (insofar as we focus on the action, rather than the agent) that though the action is of a type that is usually wrong, in these circumstances it was not wrong” (Baron, 2005, p.389). The proponent of either the Weak or Strong Difficulty Thesis claims that difficulty can be a justification condition. I, however, will argue that difficulty is of the type of extenuating circumstance that makes an action which *is* morally wrong not blameworthy: excuses.

We can make one argument for the consideration of difficulty as an excuse condition by making use of discussion from literature. First, it might be useful to think a little about the nature of excuses. Often, in everyday moral discourse, a mistaken belief can provide an excuse. Imagine the case in which an agent deliberately smashes the window on the front of someone’s house with a rock. Ordinarily we would acknowledge this behaviour as wrong, as it is wrong to damage the property of others without permission. If there is no more to the story – if this was sheer and gratuitous destruction – the

---

distribution of costs matters here. For example, we might not think that an aggregation of interpersonal costs can outweigh the cost to one individual, e.g. (Scanlon, 1998, chapter 5). Whether or how this is the case is not a matter for discussion here.

<sup>15</sup>For example, Smith (1983), Kelly (2013), Baron (2005, 2007).

<sup>16</sup>In criminal law, we can categorise these as *defences*. These acknowledge that a law was broken, but point to circumstances which make it the case that punishment is not appropriate, and (where successful) result in acquittal (Dressler, 1995, §16.01). In legal philosophy there are also two other types of defenses. These are “specialised defenses”, which are specific defenses acknowledged in the law for particular crimes, and “extrinsic defenses”, which acknowledge the culpability of the agent but offer reasons why prosecution is not appropriate (Dressler, 1995, §16.03,[D],[E]). These less obviously have parallels in the moral case.

agent is blameworthy for this conduct.

If however, the agent broke the window to gain access to the building after judging that a resident needed urgent medical attention the verdict is likely to be different. This will often provide at least an excuse for doing this, even if the agent was mistaken. Perhaps they leapt to the verdict that the resident was ill, when they were simply lying in the middle of the floor relaxing, and the agent should have knocked on the window first or done more to ascertain the particulars before responding in that manner. (If they were correct, or perhaps warranted in their judgment that help was needed, then the verdict is likely to be that the act was not even wrong, as this would justify the breaking of the window.)

So mistaken beliefs or ignorance of certain facts can absolve an agent of blameworthiness.<sup>17</sup> These types of scenarios do not usually provide an excuse in legal philosophy, as a certain *mens rea* – state of mind – is, for most crimes, a required element.<sup>18</sup> There is a structural similarity even between these cases and the moral cases however, as it is still accepted that the *actus reus* – the criminalised action – has been committed, but that the agent is not culpable.

What else can count as an excuse? There is considerable discussion of excuses in legal literature. While there are differences in what amount to excuses in legal and moral discourse, the sorts of explanation that can serve as excuses are similar. Of course, excuses that exculpate in law may not leave one without blame morally, because there is a great deal of behaviour which, though we may condemn, we do not deem suitable for criminal punishment.

Marcia Baron categorises excuses as falling into two types:

- a. those (insanity being the paradigmatic instance) that come into play because of some feature of the actor that differentiates him from most adults and makes it very difficult for him to act as the law requires.
- b. those (such as duress) that come into play because the situation was

---

<sup>17</sup>This is not always the case, as there are cases of culpable ignorance, discussed notably by Holly Smith (1983).

<sup>18</sup>There are cases where mistakes of facts (or even mistakes of law, which usually cannot provide legal excuses) are able to excuse, and arguments that further mistakes should excuse as discussed in Garvey's "When Should a Mistake of Fact Excuse?" (2009). Normally, however, defenses based on mistakes of fact will constitute "failure of proof" defences.

such that it was extremely difficult for that actor *and would be extremely difficult for most actors* to avoid acting wrongly or unlawfully.

(2007, p.23<sup>19</sup>)

It is notable that both categories reference *difficulty*. However, the descriptions here seem to relate more obviously to other sorts of difficulty. Motivational difficulty, it seems, has to be *perceived*, whereas some of the difficulty discussed here clearly does not. Its being difficult (either generally, or for a particular agent) to discern whether someone needs urgent medical attention or whether they are simply relaxing in a situation is not something the agent will notice. These are instances more like the difficulty involved hitting a bullseye, in that they are difficult because of the unlikelihood of success despite intentions oriented appropriately, rather than Cohen's case of cycling with someone on the handlebars, or the motivational difficulty previously mentioned, where the difficulty is phenomenologically salient to the agent.

Cases of excuse because of duress are helpful cases. They give an instance of non-blameworthy agents, where it seems like they are not blameworthy because of difficulty, but the obligation remains unchanged. This is evident because we still deem the agent succumbs to duress as having done something wrong (even if they are not blameworthy because of it). Duress is (usually<sup>20</sup>) accepted as a legal defense if the following holds:

1. another person threatened to kill or grievously injure the actor or a third party, particularly a near relative, unless she committed the offense;
2. the actor reasonably believed that the threat was genuine;
3. the threat was "present, imminent, and impending" at the time of the criminal act;
4. there was no reasonable escape from the threat except through compliance with the demands of the coercer; and
5. the actor was not at fault in exposing herself to the threat

---

<sup>19</sup>While Baron intends her discussion to bear fruit in legal discussions of excuses, she expects her discussion of excuses to apply to morality more generally (2005, p.39).

<sup>20</sup>Dressler notes that murder is a general exception to this (1995, §23.01[B]).

(Dressler, 1995, §23.01[B])

It seems clear that duress of this sort may excuse an agent, not only legally, but also morally. In many cases, it might look like the agent is not only excused, but also morally justified, in complying under duress. We could imagine, for instance, a legitimate threat against one's life unless one stole a snack from a convenience store. This type of behaviour, at least for consequentialists, will be entirely justifiable. When the action required under compliance is considerably worse (but still excusable), however, we may judge that the agent isn't blameworthy, but be unwilling to accept that the action was not wrong. By excusing the agent in this case, we say "Yes, you did something wrong, but considering the circumstances, we can't blame you".

It also seems like identifying difficulty as the cause for this excuse is the correct diagnosis. We can imagine an agent going through a horrible internal struggle before co-operating with a terrorist's demands and finding resisting – believing that this would result in their death – to be too difficult (maybe impossible, but for simplicity, let us not assume that at this juncture). If another agent of very cool temperament, simply evaluated the situation and decided that they would rather live succumbed to the same demands – with no sign of difficulty or internal struggle – we may judge them considerably more harshly.

In some relevant ways, the duress cases seem structurally similar to the cases involving motivational difficulty discussed earlier. The agent may realise what they *should* do, and this may involve minimal physical exertion, but psychologically it can still be extremely difficult. If we accept the symmetry in this regard, and accept (as I have suggested we should) that difficulty acts as an excuse condition in duress cases — and that the difficulty cannot, in those cases change the obligation-status of actions — we may expect a parallel verdict in other motivational difficulty cases.

Unless some relevant disanalogy can be noted, or alternative arguments provided, I suggest that this treatment is appropriate in all cases where difficulty can alleviate blameworthiness. If this is correct, then though difficulty can alleviate blameworthiness, it does not prevent difficult actions from being obligatory.

A supplementary argument for this conclusion may be made based on the phenomenology of blameworthiness for the blameworthy agent, and fe-

licitous responses for an excused agent. Consider an agent who has an obligation to  $\phi$ , but something transpires that makes  $\phi$ -ing significantly more difficult. Perhaps she had agreed to meet a friend, but her method of transport broke down and the additional difficulty this entails makes it not blameworthy to fail to  $\phi$ .<sup>21</sup> Now, we may imagine that the agent realises this about her situation — she knows that she would not be blameworthy if she does not  $\phi$ . For an agent in this situation, it would be strange to count the fact that she would be ‘off the hook’ as an allowance to not  $\phi$ . In order for the agent to not be blameworthy, it should be the *difficulty* that explains the not  $\phi$ -ing.

It would be infelicitous for the agent to count the reason that she could be excused as somehow *in favour* of not performing the required act. We might imagine an agent who has promised to meet a friend, but knows that if the weather is horrible enough, this would make motivating her action difficult enough that she would not be blameworthy for failing to do so. She stares out of the window while getting ready to leave, and when the rain picks up sufficiently says to herself with a smile “Oh good, the weather is horrible enough now, so I’m off the hook — I don’t have to go.” This somehow seems inappropriate. It would be much more appropriate for her to anxiously struggle over the decision, still perhaps feeling that she *should* go, but not being able to bring herself to do it.

An agent with appropriate sentiments will also feel differently in the aftermath both of these cases. If, as Ross discusses, an agent who deems themselves justified in breaking a promise, while realising that the prima facie reason to keep the promise remains, may feel a compunction to make amends (2002, p.28). They will not however (or, it would not be appropriate to) feel the “shame or repentance” which accompanies a wrongdoing (even when not blameworthy).

If the arguments presented in this section are successful, we should reject the claim that the difficulty of an action that otherwise would be obligatory can render it non-obligatory, i.e. we should reject the Weak Difficulty Thesis.

---

<sup>21</sup> Keeping costs fixed in this type of case will be difficult. Usually this sort of case will result in additional time, physical effort or money being expended. For the sake of the example, let us assume that costs are somehow fixed between the cases.

### 5.2.3 Resistance to Understanding Difficulty as Merely an Excuse Condition

This is however, just one type of case. We could accept that the actions of someone committing a serious crime (even under duress) will be evaluated as morally wrong regardless of the difficulty, yet say for other actions, or failures to act, the wrongness may be contingent on difficulty. That is, we may maintain that though difficulty sometimes only acts as an excuse-condition, there are instances where it can justify an action, that is, make an otherwise morally wrong action permissible.

Consider the following situation:<sup>22</sup> April is a nurse who works in a childrens wing of a hospital. April is well qualified for this job, and though it is emotionally draining, she finds it very rewarding. After an increase in crack cocaine in the area, the amount of children coming into care increases dramatically, making a tough job even harder. April realises that though she already has a job where she can help provide care for the children, if she was to decide to become a foster carer, she could take one of the children who was healthy enough to leave the hospital into her home and do even more good. She knows that while she is doing a great deal to help, she could do even better.

Generally, in this sort of situation, we would not say that the April has an *obligation* to help. It might be a really good thing to do, and extremely noble, but not obligatory. The real April actually did become a foster carer in order to take in babies and provide extra care, but it would be bizarre, if she had not volunteered to foster the children, to ask why she didn't. There are uses of the word 'ought' under which we might say that she ought act in this way, but this is not the 'ought' of moral obligation, but rather an 'optimific' ought, an ought of *most moral reason* – the type of ought that can consistently be used in phrases like “You ought to  $\phi$  but you don't have to”. It is what Fishkin refers to as a “weak ought” (1982, p.12).

Because, in this sort of situation, we would not think there is an obligation, and an obvious reason why this is the case, presumably, is just how difficult it would be for her. The apparent difficulty seems to make the behaviour optional, rather than something that she ought to do, but would be

---

<sup>22</sup>This example is modified from a similar case, based on a true story detailed in an article on the 10th of March, 2013 in the *Cincinnati Enquirer*, and also discussed by Vanessa Carbonell (2016, p.39)

excused for not doing.<sup>23</sup>

If the difficulty of an action in this sort of case is the only reason we can posit for the action's not being obligatory, and we hold firm to the intuition that this sort of action is not required, the orthodox understanding of demandingness must be reconsidered: difficulty must be attributed an integral function in our conceptions of when a theory is too demanding. This type of argument offers a convincing role for difficulty in our moral framework.

In the following section, I will attempt to demonstrate why I think this argument fails.

#### 5.2.4 Rejecting Difficulty and Reconceptualising the Problem

If we formalise this argument, I think we can discover where it fails, and from this pinpoint why examples of the type given above do not establish even the weak thesis.

- 1: There are cases where the morally best action is not obligatory.
- 2: This requires explanation – something must *make it* the case that the best action is only optional.
- 2a: The cost of actions in these cases cannot explain their optionality.
- 3: Only the difficulty of these actions can explain 1.
- 4: If (1), (2) and (3), then the difficulty of an action can (in some cases) make it non-obligatory.

Conclusion: The difficulty of an action can (in some cases) make it non-obligatory.

The first premise in this argument claims simply that morality does not require maximising. While there are 'extremists' who deny this claim

---

<sup>23</sup>The orthodox reason that might be offered for the optionality here would be the increased costs for the agent, which have obviously not been controlled for here. They might, for example, consider something like Scheffler's agent-centered prerogatives (Scheffler (1982)), but we have already noted the problems with that type of account in 4.

and minimalists (who claim that morality makes either no or negligible demands),<sup>24</sup> it is a fairly standard position. If a demandingness objection is *ever* appropriate, this premise must be true.

The second premise – which was previously only implicit – is where I will claim the serpents have already entered the garden (we’ll come to this later). The side premise [2a] perhaps does not need to be made explicit here, but I have included it as it is the denial of the orthodox position. The third premise is supported by the intuitions in the April case in the previous section. In cases like that, the difficulty seems to make it permissible not to perform the goodly act, and no other salient feature seems able to fill the explanatory role. The fourth premise seems obviously true, and, with the conjunction of the first three premises entails the conclusion.

The second premise claims that [1] requires some explanation, implying that the default requirement is that one *must* perform the morally best action. This is extremely dubious. There are many theoretical accounts which do not dictate that an agent is obligated to do *the best*. Any account which allows for the existence of supererogation – going above and beyond what is required – is committed to it not being the case that one is always required to do the best action possible. Supererogation is compatible with [2], but each case where supererogation is possible would have to be one where a special consideration applies; one that makes the best action optional. To take a mundane example, there are social niceties — offering someone a cup of tea, perhaps — which do seem good, but completely optional. A compulsively courteous host may view this as a requirement, but that will definitely not be the norm. Adopting [2] would mean that we would require a special explanation why (assuming the tea-making is the morally best behaviour) it is not a requirement. Simple mundane cases like this suggest that the presumption of a default obligation to maximise the good seems problematic.

I claim that most people who are responding to demandingness objections (other than the extremists, who simply claim that they are bad objections) embrace an argument with premises 1 and 2. This explains Scheffler’s move in adopting agent-centered prerogatives: he thinks these prerogatives explain premise 1, seemingly accepting a default position that we ought to

---

<sup>24</sup>This is again making use of Kagan’s terminology from *The Limits of Morality* (1989, p.5).

perform the morally best option. Murphy similarly endorses his collective principle of beneficence to limit our moral obligations. That he needed to do this, I claim, indicates that he assumes that moral obligations needed to be limited, that otherwise they would naturally be unlimited, and that we would be obligated to perform the best action.

So, let's think about why this presumption might hold. Why might we think that one is always morally obligated to perform the morally best action? Obviously, that best action does have the most (morally) to say in favour of it. One question we might ask is whether there is a parallel across other domains. Is there, for instance, a prudential obligation to act in the most prudential way?

Trevor Hedberg claims that there does exist a concept of epistemic supererogation (2014). Similarly, McElwee argues that it does make sense to talk about supererogation across a variety of normative domains, including prudence and etiquette (2016a). The rationale for this claim is the structural similarities in how we treat/discuss these domains. For example, there are general norms which apply to etiquette: it is expected that people wear relatively clean clothes, that they use a knife and fork in the 'proper' manner or that they begin letters by appropriately addressing the addressee. If an agent is found flouting these norms this may result in some form of social censure – they may be talked about behind their back, mocked personally or excluded from further similar social gatherings. This is the analogue of blame in the moral case. Similarly, many of these norms may be adhered to in varying degrees. The obligation of etiquette to dress appropriately for a nice dinner might be conformed to by wearing a shirt and jacket, but it might be even better to wear a suit or a tuxedo. Perhaps the person really 'making an effort' would be praised for doing so, whereas those only just adhering to the norms would receive no positive or negative attention. We might suggest that there being a requirement and the ability to do better across other domains indicates that we should think something similar occurs in the moral case.

That there are claims of supererogation in other domains, however, may not convince any advocates of [2]. They may claim that the default obligation is to do the best in all of these domains, and that some extenuating circumstances must apply in each case in order to make it permissible (in that domain) to act sub-optimally. Alternatively, they might simply claim

that morality is special. Discussions regarding the overridingness of morality might lend credence to this, and considering the arguments I have already offered (e.g. in 2.4.2, when discussing Dorsey’s position), this might seem like a plausible position.

Another way we might suggest that the thinking behind premise [2] is mistaken is again by simply considering our ordinary moral practices. Premise [2] would have us accept that moral failing – exhibiting some moral failure – normally (absent extenuating circumstances) means simply acting non-perfectly. Considering that it nobody *does* act perfectly, we must either accept that extenuating circumstances are ubiquitous or that everyone is acting wrongly practically all the time. The former of these options is clearly absurd, pathologizing the normal. The latter, however, also sounds pretty ludicrous.

What do we think our moral concepts are doing? Rather than assuming that the best is what is required, might we not suppose that obligations are grounded in other ways? Other moral asymmetries might be relevant here. We tend to consider not offering help to someone in dire straits in one’s immediate proximity to be abhorrent behaviour. However, simply not doing good things for people who are not in pain is generally considered to be morally neutral. Does this not tell us something about our concepts of moral obligation?

We might suggest that the second premise could be read so that it is only in high stakes cases – where one’s action could do a great deal of good, or prevent a great deal of bad, perhaps – where there is an expected (default) requirement to do the best. However, these are *exactly* the types of cases where we celebrate supererogation. We adulate those who run into burning buildings risking their lives to save others – this is *definitely* not ordinarily required.<sup>25</sup> We do not really care much about supererogation in a trivial case, like doing slightly more than one’s fair share of the household chores.

One way we can show that premise [2] is at least not obviously true is by considering positions which might be taken to deny it. While utilitarians are often seen as accepting something like this, Mill definitely did not,<sup>26</sup> it is not

---

<sup>25</sup>Obviously firefighters and some others will have special obligations in such cases.

<sup>26</sup>This is apparent from chapter five of *Utilitarianism*, as well as in “Auguste Comte and Positivism” (1865, p.143), where Mill explicitly endorses supererogation. See Jacobsen’s “Utilitarianism without Consequentialism” (2008) for extended exegesis on this point.

required by variants like satisficing utilitarianism,<sup>27</sup>. Rule consequentialists will also typically reject this.<sup>28</sup> From a Kantian perspective, we might embrace a position like Thomas Hill's, according to which it is permissible to discharge imperfect duties to varying extents, such that the *best* act might not be required and may instead be supererogatory (2002). We could also look towards more radical accounts, such as the type of moral bargaining endorsed by Gilbert Harman (1975; 1996). According to this view, our moral obligations depend on our agreements, so unless our agreements made it the case that the default moral obligation were maximal, it would not be, so premise [2] would be rejected.

In rejecting premise [2] for the above argument about difficulty, I need not engage with the third premise. It is strange to ask, I have argued, why we are not obligated to do the best in some situations, just as it would be strange to ask "Why are you not travelling to Madrid today?" Without some background assumptions which suggest that we have reason to expect you to be acting in this way, the question is very bizarre. I have also suggested that one reason it doesn't strike us as weird in the moral case is optimific and deontic readings of 'should' and 'ought'. This is why it seems peculiar in the case of April in the previous section to describe her as excused if she chooses not to foster the children (in addition to all the good work she is already doing). A reason why she might not perform the act is the difficulty, but she simply is not under an obligation to foster one of the children anyway, so it isn't an excuse.

If premise [2] is false, this gives us an easy response to the April case. Rather than asking, "Why isn't this required?" we ask "Well, why would it be?". The question of what *is* required will require a story about when and how obligations are generated, but so long as this does not generate what we think of as excessive requirements, it can forestall the problematic dialectic move.

In the following section, I will explore the potential of accepting something like this shift in the dialectic, and the suggestion that we instead should focus on developing an account of moral obligation that is less parasitic upon whatever theory of value we take.

---

<sup>27</sup>Slote talks about satisficing *consequentialism* more generally but notes that his discussion also applies to utilitarianism specifically (1984).

<sup>28</sup>We may consider Hooker's account from *Ideal Code, Real World* (2000) or Mulgan's from *The Demands of Consequentialism* (2001) as examples.

### 5.3 Foundations of Moral Obligation: Success Criteria for a Moral Theory

In the previous section, I suggested that a good moral theory should have a theory of obligation separate from its axiology: simply finding what acts have the most moral value does not answer all the questions about one's moral obligations. In chapters 3 and 4, I suggested that certain social factors appear to play an important role in what we take our moral obligations to be. Compliance in particular was discussed, and with it Murphy's collective principle of beneficence in chapter 4, but this was ultimately rejected in part because of its counter-intuitive implications.

Despite rejecting the account of beneficence based on deriving individual obligations from the obligations of a group, cases like Le Chambonnais or other instances where something like ratcheting-up (e.g. Carbonell, 2012) seems to occur, indicate that standards can be raised when communities accept more burdensome obligations.

Before putting forward a positive proposal, let us consider some criteria by which we may evaluate how successful a theory of moral obligation is. I tentatively suggest the following checklist of conditions a theory ought to satisfy. It should:

1. conform to most/all strong intuitions
  - a) provide 'moderate' verdicts.
  - b) support intuitions about moral progress.
  - c) must be able to satisfactorily respond to Cullity's iterative demands problem.
2. allow for moral discourse.
3. allow for the possibility of moral mistakes.

There are clearly other criteria by which one could evaluate an account of obligation, but this tentative list incorporates some of the problems that have been identified in the previous chapters. Let us consider the list briefly before considering whether it is plausible that a theory could satisfy these criteria .

## Conforming to Strong Intuitions

As, when responding to demandingness objections, we are relying upon intuitions, any resulting account had better fare well on this front.

Recently, and particularly in the last decade, a great deal of suspicion has been cast over the use of intuitions in philosophy. Some of this has followed interesting results developed through the experimental philosophy movement. As discussed in 2.3.1, Unger discussed some of the issues facing our various types of intuitions (1996), particularly raising difficulties with order effects and other biases which affect our judgments. Kagan similarly raised issues with intuitions, focusing in problems with intuitions about cases and arguing that it is likely that some form of error theory will be required (2001, p.59).

Of course, up until this point, I have been making use of intuitions. I will not offer a substantive defence of this practice here,<sup>29</sup> but will suggest that we should focus on certain strong intuitions and seek a theory that is consistent with those. Mulgan makes a distinction between *decisive* and *distinguishing* intuitions, where the former are judgments a theory *must* accommodate, and the latter are treated as secondary (2006, pp.2-3).

In chapter 2, I argued that we should adopt the moderate position, that is, that theories should make non-negligible moral demands, but not (typically) make demands as extensive as those prescribed by maximising act consequentialism. As this is entailed by the position that underpinned the demandingness objection, it will be of no surprise I regard this as a decisive intuition.

A second general intuition which would be desirable for a theory to be consistent with involves the notion of moral progress. We know that the extent of practices like slavery were far more commonplace at various times in the past, and we want to be able to say that removing these is *progress* in some way. It would be a strange and alarming result if our theory could not evaluate a situation as *better* today because of those sorts of developments.

A third feature of preserving certain strong intuitions will be the ability to resist extreme demands that result from Cullity's problem of iterative

---

<sup>29</sup>There are several ways one may defend the use of intuitions. Kagan discusses several options with these, but judges more work on the topic to be required. Perhaps the most radical treatment of the issue is Herman Cappelen's *Philosophy Without Intuitions* (2012), wherein Cappelen claims that philosophers do not actually use intuitions, but rather considered judgments with arguments.

demands. As I suggested that some accounts (like Scheffler's agent-centred prerogatives) should be rejected because of an inability to respond adequately to these cases, a theory that is able to give moderate intuitions in these cases is required.

### **Moral Discourse and Moral Mistakes**

Conforming to intuitions about how demanding morality is seems to be the most obviously relevant criterion of evaluation, when a previous theory has been rejected for being too demanding. In addition to this, I have suggested that a good moral theory should permit moral discourse and should be compatible with the notion of moral mistakes. A major motivation for this will become apparent in the next chapter, as these are often seen as providing problems for the type of account I defend. With that in mind, I will say little further on this now, but take it for granted that it is clear that these are desirable features of a moral theory.

In this chapter, I have analysed the notion of blame and suggested that the implication of blame is a major motivation in making demandingness objections persuasive. Through rejecting McElwee's position on difficulty, and considering when excuses are appropriate, I suggested that some of the discussions in this area seem predicated upon an assumption that one has an obligation to do perform the best action. I have also suggested that considering how obligations are formed, and being sensitive to ways this may incorporate social factors, may help us respond to overdemandingness. In the final chapter, I offer a schematic for how this could be done.

## 6

# Positive Proposal: Reasonableness in Obligation

*“Kimmy, we have a wonderful thing here in New York called ‘moral relativism’.”*

– Lillian Kaushtupper, *The Unbreakable Kimmy Schmidt*<sup>1</sup>

In this chapter, I offer a tentative proposal of a theory of moral obligation which can provide a response to the demandingness objection, and satisfy the success-conditions laid out at the end of the previous chapter.

In chapter 3, we noted that several traditional explanations for what makes a theory overly demanding seem to fail. It was also suggested that some features – like whether others are in a similar position to act – which are often taken not to be morally relevant do seem to make a difference to our intuitions about what obligations are appropriate. Singer raises this explicitly in “Famine, Affluence and Morality”, claiming that both in the case of the drowning child and the Bengali refugees, whether there are others in the same position “can make no real difference to our moral obligations” (1972, p.233).

It is puzzling, given Singer’s claims, that our moral intuitions about what is obligatory *do* seem to be affected by what others expect and what others like us are doing. In the preceding chapter I explored whether the notion of difficulty might offer a successful explanation for this. The difficulty of an action appeared to be a promising candidate because it could depend in part

---

<sup>1</sup>Season 2, episode 1: “Kimmy Goes Roller Skating”. Written by Tina Fey.

on social factors, and perhaps in turn this could limit our obligations. In 3.3, we noted that ‘social difficulty’ could explain some of the phenomenology. Seeing many others act in a certain way may make it appear easier, which could remove actual psychological barriers. If, on the other hand, no one else is acting in some self-sacrificial way, or if no one expects people to act in that way, it can make it difficult for an agent to act in that way. Similarly, if an agent *is* expected to act a certain way, social pressures to act in that way could make it more difficult to *not* act that way. Given this, it seemed plausible that as theories typically do not account for difficulty, this could provide an explanation for why so many seem to demand so much in certain situations.

The particular reason difficulty did not provide a suitable explanation for overdemandingness was that difficulty is an *excuse*, not a justification. This is important because when excuses are applicable the conduct was still morally wrong – i.e. the agent was still under an obligation not to act as they have (not  $\phi$ -ing was still *demanded* of them). In order for not accounting for difficulty to explain why a theory is overly demanding, an act’s difficulty would have to be a *justification* for not doing it, i.e. it would have to make not performing that act permissible.

The more general diagnosis that emerged is that it is a mistake to try to explain why obligations could vary with these social factors (and relatedly overdemandingness intuitions) by looking at how our obligations might be *limited*. I proposed that instead, it might be more promising to look at how the obligations are generated. This way of thinking does not suggest that, as a default, optimal acts are obligatory.

The problem then arises again of how we might explain how obligations can vary with social factors, if not because of difficulty.

Let us consider situations where intuitions about our obligations seem to vary in accordance with one’s moral community. For example, it seems clear that in a community where the vast majority of people recycle their plastics and expect compliance with this practice, people within that community are more likely to judge that behaviour to be obligatory compared to a community where compliance with and expectation of this practice is negligible.<sup>2</sup> Why might the members of this moral community judge that

---

<sup>2</sup>This example is structurally the same as the Chambonnais case, discussed in section 3.3, but is perhaps a more commonplace type of case.

they were not under an obligation to recycle here? In such cases, it might even be acknowledged that one rule – like recycling one’s plastics – would be ideal, but agents do not consider it obligatory to act optimally. This case seems parallel to the situation of many of us with regards to the global poor. Many people will acknowledge that the state of affairs would be vastly improved if we all donated in line with a fairly burdensome rule, but not deem it obligatory. Cases like these may prove elucidating, as the intuitions involved in them form the bedrock for making a demandingness objection. If we judge that these intuitions are veridical – which it seems we must if we accept (as I have suggested that we should) that the demandingness objection offers a genuine challenge – we require a moral theory that can allow for obligation-verdicts to differ because of these factors.

This chapter provides a proposal of how demands might be generated in those situations. I attempt to suggest that by making use of reasonableness – a concept which already features heavily in moral theory – we could explain why social factors play a part in generating obligations.

We can see that reasonableness appears important from looking at some descriptions of the problem of overdemandingness. Consider, for example, Tim Mulgan’s description of The Demandingness Objection:

“This leads to the common objection that Consequentialism is unreasonably demanding, as it leaves the agent too little room (time, resources, energy) for her own projects or interests. I shall call this the Demandingness Objection.

(2001, p.4)

And another description of the objection by Hooker:

“The objection levelled at act-utilitarianism is that act-utilitarianism demands more sacrifice than it is reasonable to demand”

(2009a, p.151)

A third example comes in Liam Murphy’s discussions, where he claims that when . . .

“. . . we ask what an acceptable level of demands would be, all we can say is that the demands made by the morality we are committed to are reasonable.”

(2000, p.68)

And here is a fourth from Fiona Woollard:

“The Demandingness Objection is the objection that a moral theory or principle is unacceptable because it is too demanding: it asks more of an agent than we can reasonably expect.”

(2016a, p.89)

A common feature in these descriptions is an allusion to what is reasonable.<sup>3</sup>

We might note ambiguity present in each of these quotes. It is not clear from these descriptions whether reasonableness is a *constraint* or *limitation* on demands, or constitutive somehow of what they could be. If anything, a more natural reading would suggest that reasonableness is only a limitation on moral demands. An alternative way that we might view reasonableness is as actually having a role in generating obligations or demands, rather than merely *limiting* them so that they are more modest than the requirement to act optimally. Later in this chapter an account of this sort, provided by T. M. Scanlon is discussed. His account sees what is reasonable as having a role in determining what the demands are.

It is claimed that moral demands *must be reasonable*. Why might this be the case, and what is the content of this requirement? We can also note that several accounts already give reasonableness a central role, thus we might speculate that we could provide an adequate account of moral demands such that it had at its heart a concept of reasonableness.

A major issue for some theories seems to be that our moral intuitions are affected by social factors (expectation and compliance). Why might an account which appropriately regards reasonableness be any more able

---

<sup>3</sup>This type of description of the problem - alluding to the failure of a theory to make reasonable demands, or being unreasonable in its prescriptions - is adopted by several others too, including Ashford (2003, p.273), Cullity (e.g. 2004, p.101), Chappell (2007), Yetter Chappell (2017).

McElwee talks of the objection in terms of what it would be reasonable to blame people for, claiming an overly demanding theory makes prescriptions that it would be inappropriate to blame people for feeling to meet (2016b, p.19), and elsewhere of theories providing “verdicts that are unreasonably demanding” (2017, p.84). Woollard gives a similar description in “Dimensions of Demandingness”, noting that “some theories or principles ask so much of the agent that they mean it is not reasonable to expect an agent to choose to conform to the moral principle” (2016a).

to respond to this problem? This shall be developed further later,<sup>4</sup> but to provide a brief response, we can see reasonableness as being sensitive to these issues. If it is expected of everyone in a community that they will recycle their plastics, not doing so is unreasonable. What and how much it is reasonable to do in a given community will vary with what people *actually* do. In this way reasonableness may have a constitutive role in determining the obligations an agent has, and – or so I will suggest later – may evade difficulties posed by the some of the intuitions other moral theories (like maximising utilitarianism) struggle with.

Before seriously analysing the notion of reasonableness, a brief clarificatory note should be made regarding the meaning of the term throughout history. At least at some points in the past “reasonable” simply seemed to mean *rational* or *in accordance with reason*, whereas “unreasonable” meant *contrary to reason*. We can see this usage in Hume’s *Treatise on Human Nature*:

“It is not contrary to reason to prefer the destruction of the whole world to the scratching of my finger. . . . In short, a passion must be accompanied with some false judgment, in order to its being unreasonable; and even then it is not the passion, properly speaking, which is unreasonable, but the judgment.”

(Hume, 1969, II.3.3)

Hume simply equates the unreasonable with that which is contrary to reason. The meaning of “reasonable” has clearly changed since Hume’s time. In today’s parlance it definitely *would* seem unreasonable to prefer the destruction of the world to a negligible harm, though perhaps not *irrational* if the agent had very strange goals or desires. This distinction between *reasonableness* and *rationality* is the focus of some discussion, notably by Scanlon.<sup>5</sup> What the current meaning of the term is – and the concept that I suggest is useful in successfully accounting for our moral demands – will be discussed in the following sections.

Putting reasonableness at the centre of moral theory is not a new idea. In *What We Owe To Each Other*, Scanlon writes that he takes the notion

---

<sup>4</sup>See section 6.1.5

<sup>5</sup>This is discussed further in section 6.1.3.

of “reasonable” to be “basic to moral thinking” (1998, p.192), and it plays a significant role in his contractualism via reasonable rejection. Margaret Moore notes that as well as Scanlon, Brian Barry (1989) and John Rawls’s (1991; 1993) accounts also put reasonableness at the forefront of moral theory (1996, p.168).

In the following section, I examine several discussions of reasonableness in the last century, highlighting the key features identified. In section 6.3, I discuss Scanlon’s account, as an example of an extant theory which heavily relies upon reasonableness, before suggesting reasons why it might be overly demanding, offering a challenge broadly related to Margaret Moore’s criticisms. I then tentatively offer an alternative account of what reasonableness entails, and drawing upon Gilbert Harman’s moral relativism, suggest that such an account could prove successful in responding to many of the problems raised in previous chapters. Finally, I evaluate how successful a theory of this sort might be in view of the issues raised for other accounts in chapter 2, the success criteria discussed in section 5.3 and objections unique to a relativistic accounts of obligation.

## 6.1 Views of Reasonableness

### 6.1.1 Sibley and Rescher

Serious discussion of reasonableness occurs in two articles from the 1950s. Both focus on distinguishing rationality from reasonableness. We can see that this is an important distinction to make if we consider the variety of theories that make use of one or other of these as central concepts. As Scanlon notes (1998, p.190), each of Gauthier, Hare, Kant, Habermas discuss rules (or principles) that everyone would (or should) rationally agree to.<sup>6</sup>

Sibley complains that many writers “assume that these terms [‘reasonable’ and ‘rational’] (and their antonyms) are completely synonymous” (1953, p.554). Sibley argues that this is not the case. In doing so, he suggests that rationality is an “intellectual virtue” and that reasonableness is instead a “moral virtue” (1953, pp.556-557). More specifically, he claims that reasonableness requires “a positive sympathetic disposition towards

---

<sup>6</sup>Scanlon then attempts to demonstrate that reasonableness is a better candidate for a source of normativity. This is discussed in the next subsection, and more extensively in 6.3.

others” (1953, p.557).

Rescher largely agrees with Sibley, adding that the distinction is between “reason and judgment, rationality and reasonableness, shrewdness and wisdom” (1954, p.61). Acting reasonably, however, seems for both Rescher and Sibley to require rationality. Rescher states that “whereas any discussion presupposed the rationality of the interlocutors, we must add in discussing ethics the tacit premise that they are reasonable” (1954, p.61).

The general picture that both endorse seems to be that rationality requires an ability to reason adequately, but does not require that an agent possess any specific values. Returning to the Hume quotation from earlier, we would not claim that there is anything *irrational* about preferring the destruction of the world to a scratch on one’s finger,<sup>7</sup> but it would not be *reasonable*. Being reasonable requires accepting certain things as valuable. Sibley describes this by saying:

“If I seek, for example, to justify my action by pointing to some good it achieves for me, I must be prepared to allow, as at least a prima-facie objection to it, that it results in what another deems harm to himself. Reasonableness thus requires impartiality, “objectivity”; it expresses itself in the notion of equity.”

(Sibley, 1953, p.557)

Sibley claims that in order for an agent to be reasonable, they must be able to justify their behaviour by use of reasons they would accept from others. The onus on *impartiality* in reasonableness is discussed in more depth later.

### 6.1.2 Rawls

Sibley and Rescher clearly saw reasonableness as an important moral notion. Sibley went as far as stating of someone who is merely rational (not reasonable) that “we do not know what ends he will aim at in his conduct; we know only that, whatever they are, he will use intelligence in pursuing

---

<sup>7</sup>Of course, if one does hold some values that are better promoted better by the latter than the former – as even the most thorough-going egoist usually would – it would be irrational too. The claim here is simply that there is nothing *necessarily* irrational about having that preference.

them”, whereas of a reasonable person we may “infer that he is willing to govern his conduct by a principle of equity, from which he and they can reason in common” (1953, p.560). But while they valued its importance, contract theorists like Rawls and Scanlon made further use of the notion, placing it at the forefront of their moral thought.

In *Theory of Justice*, rationality receives more of the focus than reasonableness. It is ‘mutually disinterested rationality’ that agents should possess in the original position (1991, p.13, et passim). He later conceded, however, that there is a pluralism of “incompatible but reasonable comprehensive doctrines” (1993, p.xviii). As agents holding these incompatible doctrines will have conflicting views, it is a challenge for political liberalism to demonstrate how people with these incompatible views might coexist and what must happen in cases of conflicts. A further issue, and at issue here, is how one determines which of these comprehensive doctrines are reasonable.

Like Sibley and Rescher, Rawls distinguishes the reasonable from the rational.<sup>8</sup> Rawls claims that reasonable people. . .

“... are not moved by the general good as such but desire for its own sake a social world in which they, as free and equal, can cooperate with others on terms all can accept. They insist that reciprocity should hold within that world so that each benefits along with others.”

(1993, p.50)

Once again, an emphasis on reciprocity can be noted. As a hallmark of reasonableness, reasons that are deemed as appropriate for one person must also be taken as appropriate for others.

For Rawls, justice as fairness requires that an agent must be both rational *and* reasonable. He does seem to accept that agents can actually be (at least in theory) *merely* reasonable, as well as *merely* rational. Merely reasonable agents “would have no ends of their own they wanted to advance by fair cooperation”, whereas merely rational ones would “lack a sense of justice and fail to recognize the independent validity of the claims of others” (1993,

---

<sup>8</sup>Rawls actually traces this distinction back to Kant, claiming that the rational/reasonable distinction is analogous to hypothetical/Categorical imperatives (1993, p.48, n.1). As will be seen later in this chapter, I do not think the distinction maps quite so neatly.

p.52). From this it seems we can pinpoint what reasonableness entails for Rawls: a sense of justice and the ability to recognise the independent validity of claims of others.

He later elaborates that the reasonable is “not the altruistic” (1993, p.54); it simply requires reciprocal acceptance of reasons. The reasonable agent will treat a reason as just as morally relevant whether it is their reason or someone else’s. The account detailed in *Political Liberalism* then requires that the principles, in order to be acceptable in a well-ordered society, must be acceptable by agents who are reasonable and rational. He summarises what the reasonable entails by discussing what a reasonable society looks like:

“In a reasonable society, most simply illustrated in a society of equals in basic matters, all have their own rational ends they hope to advance, and all stand ready to propose fair terms that others may reasonably be expected to accept, so that all may benefit and improve on what every one can do on their own. This reasonable society is neither a society of saints nor a society of the self-centered. It is very much a part of our ordinary human world, not a world we think of much virtue, until we find ourselves without it. Yet the moral power that underlies the capacity to propose, or to endorse, and then to be moved to act from fair terms of cooperation for their own sake is an essential social virtue all the same.”

(Rawls, 1993, p.54)

This captures several important features of reasonableness. It does not require perfect behaviour. It requires acceptance that others have rational ends that they want to pursue, and that these count equally when we propose or endorse moral rules. Rawls regards the willingness to propose and endorse terms of cooperation as the basic aspects of the reasonable.

### 6.1.3 Scanlon

Scanlon’s contractualism is the leading moral theory which regards reasonableness as having a crucial role in determining which acts are wrong, via his notion of “reasonable rejection”. Unlike Rawls, who permits principles

that only rational and reasonable people *would accept*, Scanlon claims that “an act is wrong if it would be disallowed by any *principle* that no one could reasonably reject” (1998, p.197). This is a further elaboration of Scanlon’s position that justification – which requires reasonableness – provides “both the normative basis of morality of the right and wrong and the most general characterization of its content” (1998, p.189).

Like Sibley and Rawls, Scanlon distinguishes the rational from the reasonable. He provides an example wherein parties are negotiating for water rights, but one landowner has control over the vast majority of the water already. In this instance, Scanlon suggests it would be reasonable for him to make an agreement whereby everyone got at very least some minimal amount of water. Not to do so would be unreasonable, however, it might not be irrational. What would be rational for him would “depend what his aims are” so presumably concerns some sort of means-ends reasoning (1998, p.193).

Like Rawls, Scanlon takes reasonableness to include some sort of recognition of the reasons/needs of others. He suggests that a reasonable judgment is “not a judgment about what would be most likely to advance their interests or to produce agreement in their actual circumstances or in any more idealized situation, but rather a judgment about the suitability of certain principles to serve as the basis of mutual recognition and accommodation” (1998, p.194).

Both Scanlon and Rawls hold that reasonableness commits an agent to some sort of values. The reasonable agent must recognise some goods valued by others and acknowledge their pursuit of those goods. This acknowledgement, if reasonable, must not discriminate based on who possesses the reason; a reason ‘counts’ the same whether it is held by them, one of their loved ones or neighbours, or a stranger. In this sense, their view of reasonableness requires impartiality.<sup>9</sup>

---

<sup>9</sup>It is notable that though this is a requirement for impartiality, what this requires is not something that obviously leads to extreme demands. In terms of the notions of impartiality Hooker speaks of (discussed in 3.2) this may require impartial application of moral rules and (perhaps) impartial assessment of first-order moral rules, but it certainly does not require impartial benevolence. Recognising that people have reasons that ‘count’ the same as one’s own does not entail that any of those reasons ground obligations.

#### 6.1.4 Marcia Baron on Reasonableness

Marcia Baron has recently written about the notion of the reasonable in both ordinary and legal contexts (2011; 2017). Baron expresses concern about conflation of the term in the legal setting. Rather than confusing reasonableness with rationality, as was Sibley's worry, in law it is "often treated as synonymous with... 'average', 'ordinary', 'law-abiding', and 'prudent'" (2017). There are many dangers, as well as clear conceptual faults of identifying "reasonable" as "ordinary". Holding those of low intelligence or children to "ordinary" standards of what they should be aware of would be a clear conceptual mistake.<sup>10</sup> In our everyday notion of reasonableness, however, we can see that a person of low intelligence might act reasonably (or at very least, not *unreasonably*), but still not be aware of factors that would be salient to others. An appropriate understanding of reasonableness must be able to account for this.

Baron's analysis of the concept is somewhat broader than that of Rawls and Scanlon. The concept contains some notion of what is *appropriate*. Let us consider the following remark on the concept:

"Reasonableness is primarily about how one reacts to others and, more broadly, how one views others. It calls for fairness, and for awareness of others as people like you with aims, interests, needs, the right to make claims on others, and bad days. It calls for tolerance and some modicum of cooperation and reciprocity. It does not call for moderation in general, but does call for moderation insofar as the moderation is in order for getting along with others."

(Baron, 2017, pp.24-25)

She notes that what is relevant for establishing an agent's reasonableness is their reactive attitude to some situation (2017). A reasonable person is not furious over minor inconveniences, does not become profoundly aggrieved by a perceived minor slight and does not blame those who do not deserve it.

---

<sup>10</sup>This type of mistake in using the reasonable person standard was a problematic mainstay in English law, as demonstrated in several high-profile cases, notably *Elliott v C* ([1983] 1, W.L.R. 939) and *R v G* ([2003] UKHL 50), where children were originally found guilty of arson because they were not as aware as we would expect from adults of damage that would result from their behaviour.

Importantly, a reasonable person would not act upon such reactive attitudes. This stage in between reaction and action gives an agent the opportunity to act reasonably. The person could then be described as unreasonably angry, but may still *act* reasonably. A hot-tempered person, with the knowledge that he is often disproportionately angered by small matters may decide to take a deep breath and calm himself. Alternatively, he may fuel his anger and cause himself to lash out. When there is no good reason to do so, “stoking one’s anger” seems unreasonable (2017, p.11).

Here we can see that an agent’s being reasonable will have an *internal* element. An agent who deliberately allows herself, or even revels in feelings of anger or cruelty may be being unreasonable.

In the requirement here of viewing others appropriately, we may be reminded of the human perspective (contrasted with the objective attitude) discussed by Strawson (2008, p.9). We may also note that perfect moderation is not required, but the amount needed for *getting along with others* is.

Baron notes that though the concepts are different, what conduct is considered “ordinary” or “customary” can *inform* what it is reasonable to expect of an agent. She considers a surgeon charged with negligent homicide (2011, p.28). In a legal case, when discerning whether or not a defendant was negligent, it might be useful to know what procedures are customary, or expected in that situation. Knowledge of the accepted practice can inform what sort of behaviour deviates from this, and whether a surgeon might reasonably have made the mistakes the defendant has.

As well as comparing reasonableness to standards like “averageness” or “law-abidingness”, in attempting to discover what type of property reasonableness is, Baron compares it to various virtues. In some respects, reasonableness shares a lot with typical virtues. With practice a person can become more reasonable, just as one can work on becoming more generous or more brave. Being reasonable is also a valuable character trait. However, reasonableness “is not exactly an excellence; it is more a threshold notion. That said, one can excel at being reasonable.”

To illustrate what reasonableness does not entail, Baron compares it with phronesis, the virtue of practical wisdom. While phronesis is a genuine excellence, reasonableness is not; it is a threshold skill – and one with “not a very lofty” threshold at that (2017, p.8). In addition:

“Phronesis involves getting it right. Reasonableness, even supreme reasonableness, does not. The phronimos cannot be a bungler (though it is consistent with being a phronimos that one occasionally gets it wrong); but a reasonable person, even a supremely reasonable person, can, and in more than one way. Even a supremely reasonable person might have a tendency to be forgetful, or to send e-mails with egregious typos, or to hit ‘reply to all’ when this was not what she intended and is a very unfortunate mistake.”

(Baron, 2017, p.8)

This is important to note. Reasonableness does not require never making mistakes. This ties partly in with the distinction between reasonableness and rationality discussed by Sibley, Scanlon and others, as being reasonable (even supremely so) does not require flawless rationality. Also, not being an excellence, but a threshold notion, it does not require optimal performance on every occasion. A reasonable action might be ‘good enough’ without being the best action available.

Several of the features of reasonableness Baron picks out are not discussed at all by Sibley, Rawls or Scanlon. We can make use of some of this evaluation in considering a more thorough interpretation of reasonableness. Of particular note for the tentative account I offer is the suggestion that what is reasonable is *informed* by the community one is in – the “ordinary” or “customary”. If we accept that these matters of convention – these contingent features of one’s environment – play a role in determining what is reasonable, and suggest that reasonableness is a constitutive factor in obligation generation, this might be used to explain how social factors can affect what morality demands. This point might be accepted even if the degree to which customs “inform” what it is reasonable to expect is fairly minimal, as Baron seems to hold. We might (as I will suggest later) postulate that these effects are actually very significant indeed, and thus have a considerable impact on our obligations.

### 6.1.5 General features of reasonableness

Here we are interested in reasonableness as it pertains to the moral domain. However, more generally the concept can be applied to all sorts of things. An

explanation may be reasonable. A theory could be reasonable. A computer's processing speed might be reasonably quick or slow. The price of a pint of cider may be reasonable. When we take these descriptions to be appropriate, we signal that whatever described is somehow acceptable (or unacceptable) – we think it conforms (or fails to conform) to the norms of the domain of discourse.

More specifically, even moral reasonableness seems to be a broader notion than some authors have indicated. As well as individual acts, a person who regularly acts reasonably is a reasonable person. It is evident that it is used in very different ways in the discussion of how ethicists have used it: reasonable *actions*, how reasonable people *act*, reasonable *rejection of principles*, reasonable *reactions* to the actions of others. Any of these could conceivably play a role in a theory of moral obligation.

As noted in the case of feeling angry unreasonably, the concept can apply to actions or internal mental processes. Reasonableness also seems to concern how one judges others. It would be unreasonable to be overly harsh in one's evaluations of others, or expect too much of them. Evaluating actions (or internal mental processes) is what we are concerned with when making moral judgments.<sup>11</sup>

In accepting that intra-mental processes relate to reasonableness, in a reasonableness-based theory of obligation we can find support for many of our common practices. If an agent has the ability to rein in their emotions, but chooses not to, we deem them responsible for any subsequent actions, even if later, overcome by emotions, they lose control. Accepting that the nature of specific internal processes of this nature may confer reasonableness upon an action, this might lend itself to a theory of action that concerns attempts. In relation the discussion in section 5.2.2 of agents who attempt some wrongdoing, but fail due to unexpected circumstances, understanding wrongdoing in terms of reasonableness would allow us to maintain that the agent still does something wrong. This would allow us to preserve the common-sense claim that people are only blameworthy if they do something wrong, which is difficult to maintain on some accounts (e.g. those which require some negative consequence to identify an instance of wrongdoing).

---

<sup>11</sup>This seems to be the case even when the moral judgments are, on the face of it, of moral principles. When we evaluate the principle "gratuitous torture is morally wrong", for instance, we are concerned with *actions* of gratuitous torture.

Reasonableness requires that when moral agents are engaged with, they are regarded *qua agents*, and not simply as things. This is the case insofar as the types of things that an agent recognises as being valuable for them must also be recognised as important to others. Reasonableness thus requires the type of weak impartiality discussed by Scanlon and Rawls – that what an agent regards as ‘counting’ morally must be regarded as counting for everyone. This does not require (at least without other factors being involved) that any of these values ground obligation upon other moral agents, but if they do (and if they do *in virtue of being objects of desire*), then they should ground obligations agent-neutrally.

It also seems correct that what actions count as reasonable – and what morality requires of agents – will depend in part on what is customary or ordinary. To give an everyday example of this, a colleague recently posted on social media asking advice, about renting practices in Edinburgh. She asked whether it was reasonable to be asked to pay six months deposit before moving in (the responses, resoundingly, said it was not). The question was clearly intended to ascertain whether this practice was typical – whether this was customary.

Because what customs are in place or what is considered ordinary is a contingent matter – the customs in place could have been different – this introduces an element of context sensitivity to the notion. This feature will become important when providing the blueprints of a reasonableness-based theory of obligation, as it can grant an escape from overly demanding moral obligations in some circumstances. Such an account thus is able to evade demandingness objections, or so I will suggest in section 6.4.

## 6.2 Reasonableness in Theory: A Very Brief Sketch

Before going any further, we can consider what might be learned up to this point. If my hypothesis is correct, the notion of reasonableness is partly constitutive of what can be morally required. Rather than viewing reasonableness as merely *limiting* the obligations of agents, it plays an earlier conceptual role: we do not examine what (absent some sort of defeaters) ought to be done and then ask whether this is reasonable (perhaps reducing the burdens accordingly). Instead, we ask what it would be reasonable to expect of an agent as a theoretically prior question.

As previously mentioned, I do not want to commit myself to a particular axiology here. Insofar as this is possible, the following should remain neutral regarding questions of the nature of value. What is desired is that the resultant theory is able to provide consistent (non-contradictory) judgments about which actions (or perhaps behaviours or dispositions) are permissible or impermissible. We also want the theory's verdicts on these matters to bear some similarity to those we ordinarily make. Perhaps this need not require perfect matching with an individual's intuitions,<sup>12</sup> but if an aim of this move is to avoid counterintuitively demanding verdicts, the theory must at least be able to typically avoid what are seen as *clearly* overly demanding requirements.

What is desired is, simply put, something to complete the sentence:

1.  $\phi$ -ing is impermissible if and only if . . .

where  $\phi$ -ing is an object of moral judgments.

And do so in such a way that largely conforms to our demandingness intuitions.<sup>13</sup>

By making use of the deontic relations, completing those conditions can give us all the traditional normative categories – permissibility, impermissibility, obligatoriness, optionality and omissibility.<sup>14</sup>

Even given everything said so far, there are many ways one might try to complete the impermissibility condition. Consider:

$\phi$ -ing is impermissible if and only if

1. it is reasonable to require an agent not to  $\phi$
2. a reasonable principle forbids  $\phi$ -ing.
3. a reasonable person would not  $\phi$ .

---

<sup>12</sup>That might not be possible without resorting to a crude form of intuitionism, whereby what is permissible is relative to different individuals at different times and is simply identical with what they view as impermissible. Peter Unger's demonstration of the susceptibility of our intuitions to order effects and various other biases (discussed in section 2.3.1) suggests that the intuitions regarding a given case is taken to be good evidence (by most of us) that this faculty is *at least* not infallible.

<sup>13</sup>Phrasing the requirements this way does not seem to load the question unnecessarily, as any of the main rivals of varying types of moral theory can provide the conditions where some  $\phi$  is permissible/impermissible, with the exception of those which do not involve obligations (but we excluded those from this general discussion in section 3.1).

<sup>14</sup>The relations are described formally in (McNamara 2014).

4. any principles which no one could reasonably reject forbid  $\phi$ -ing.
5. no principles which no one could reasonably reject permit  $\phi$ -ing.
6. all reasonable principles forbid  $\phi$ -ing.
7. a reasonable judge would condemn  $\phi$ -ing.
8.  $\phi$ -ing is unreasonable.

Some of these candidates might be plausible. Some may be equivalent (at least under some interpretations). There are certainly other ways one might compete the condition making use of reasonableness. I will not attempt to rule out all of these (or the multitude of other) potential cases. However, I will attempt to avoid the question of whether the condition focusses on principles/dispositions or individual acts – whether the  $\phi$  should refer to a principle or an action. If defending a theory, like rule-consequentialist accounts, which refer to principles, I suggest that translating my discussion into their vocabulary would be relatively simple. As Scanlon notes:

“To justify an action to others is to offer reasons supporting it and to claim that they are sufficient to defeat any objections that others may have. To do this, however, is to also defend a principle, namely one that claiming that such reasons are sufficient grounds for so acting under the prevailing conditions.”

(1998, p.197)

Moral particularists, like Dancy (e.g. 2004), will disagree with this characterisation, so a depiction of conditions of rightness in terms of actions might be more acceptable to people of that ilk. However, if we accepted something like the Kantian notion of imperfect duties, referring to individual acts might be problematic, because at one particular time an action might not be required, but it might be flouting one’s duties never to act in that way. Regardless, I do not intend to take a commitment to either position on this matter here.

That said, I propose a condition similar to the first of those listed above. As a starting point, we know from the various formulations of the objection that we want morality’s demands to be reasonable. Assuming reasonableness

is constitutive of moral obligation (as suggested so far), we might suggest something like:

**$\phi$ -ing is impermissible if and only if** it is reasonable to require an agent  
not to  $\phi$   
where  $\phi$ -ing is an object of moral judgments.

At this juncture we may note some similarity to Scanlon's account of wrong action. Scanlon makes use of principles rather than actions, and talks of principles "no one could reasonably reject" disallowing an action, rather than simply it being reasonable to require agents not to perform that action. However, depending on the interpretations of various terms, these could be equivalent. The following section examines Scanlon's account more thoroughly, ultimately noting significant points of disagreement.

### 6.3 Reasonableness in Theory: Scanlon's Contractualism

In the section 6.1.3, Scanlon's contractualism was discussed. Given the prominent role it provides reasonableness, it might seem like a problem for an account of the sort I have proposed – given that a main aim is to avoid demandingness objections – if Scanlon's contractualism is also an overly demanding theory. I contend that what leads Scanlon to his overly demanding theory is his particular conception of reasonableness; in particular, what he deems required for reasonable rejection.

In this section, I discuss how Scanlon's contractualism arrives at moral verdicts, examining the role of reasonableness and how Scanlon's understanding of the notion gives him moral principles which (he thinks) are justified. I then expose some issues with this method, suggesting that the assumption of an objectively privileged standard of reasonableness is unjustified and that it is the source of overdemandingness is Scanlon.

#### 6.3.1 How Demanding is Scanlon's Contractualism?

Elizabeth Ashford (2003) has argued that contractualist theories appear extremely demanding. In assessing how demanding Scanlon's account is, we must examine how verdicts are formed in the theory.

Beginning with the claim that what we can justify to others determines our judgements of right and wrong (1998, p.4), Scanlon arrives at the claim that an “act is wrong if its performance under the circumstances would be disallowed by any set of principles for the general regulation of behaviour that no one could reasonably reject as a basis for informed, unforced, general agreement” (1998, p.153).<sup>15</sup>

It might seem at first glance that this type of theory would be promising in providing a relatively undemanding account, because there are so many reasons that agents might reject a principle, including how demanding it is. Ashford herself notes that Scanlon’s account had a “particularly strong prospect” of providing an account inkeeping with our intuitions (2003, p.274). A principle may be reasonably rejected if an agent has a reason for rejection which has “justifiability to others on grounds they could not reasonably reject” (1998, p.218).

The reason must be held by some particular individual – an objection is not made stronger by being a reason for several individuals. Scanlon elucidates this with the case of Jones, an worker at a television broadcasting station (1998, p.235). Some electrical equipment falls on Jones during a World Cup match, causing him extreme pain. The equipment could be moved, but doing so would require stopping the transmission for fifteen minutes, possibly causing frustration to hundreds of millions of people. A simple utilitarian theory would suggest we should leave Jones to suffer until after the match is over because the aggregate enjoyment of the viewers (if there enough of them) could outweigh Jones’s suffering. Because of Scanlon’s constraint that objections be personal, he is able to give the verdict that regardless of the amount of people watching, that Jones should be saved immediately. Jones has a strong reason (based on his intense suffering) to reject any principle that does not lead to him being saved immediately. Each viewer has an objection to principles that lead to Jones being saved immediately – because their amusement would be stifled – but the reason of each is outweighed by Jones’s objection, so the viewers do not have a reasonable rejection.

Reasonable rejection can be on many grounds, and while Scanlon does

---

<sup>15</sup>Scanlon’s discussion in *What We Owe to Each Other* is largely limited to a narrow ‘inter-personal morality’, i.e. obligations that people have towards each other. This does not include things like obligations to animals, plants or artworks (1998, p.180), though he does acknowledge that those things could be wronged in the broad sense of morality.

not provide a list of potential grounds for rejection, a principle's demandingness seems like a plausible candidate. He does say, for instance, in relation to principles requiring giving one's own interests no additional weight than those of others, that one ground of reasonable rejection could be a person's "not wanting to be bound, in general, by such a strict requirement" (1998, p.225). The legitimacy of reasons like this under Scanlon's account seem to make room for a denial of demanding principles.

However, though it is not the only reason, often "the gains and losses in well-being (relief from suffering, for example) are clearly the most relevant factors in determining whether a principle could or could not be reasonably rejected" (1998, p.215). As any burdens must be deemed to be acceptable to those who would bear them, Scanlon himself notes that "the principles that meet this test may still be very demanding, and they go as far in recognizing the claims of each as is compatible with similar recognition of the claims of all" (1998, p.171). If we consider principles of beneficence, it seems that principles which make it optional to give to effective charities (in situations like those that affluent westerners find themselves in today) will be extremely burdensome on those who would not receive the benefits those charities would provide – they would suffer the burdens of disease/starvation and perhaps premature death.

Ashford actually suggests that contractualism may fare even worse than consequentialism (in terms of demandingness) as the extreme demands seem to persist in any "practically realizable state of the world" (2003, p.294).<sup>16</sup> In addition to rejecting many principles which would be less demanding for us on the grounds that they would give a stronger objection to members of the global poor, Ashford suggests that Scanlon's contractualism also would prohibit things that consequentialism typically would not, like personal air travel. This is because principles permitting actions which have a chance – even a very remote one – of being harmful may be reasonably rejected by agents who cannot benefit (or expect to benefit) from those actions (2003, p.208).<sup>17</sup> Whether or not contractualism is actually in a *worse* position with regards to the objection, if we think consequentialism is overly demanding,

---

<sup>16</sup>This claim is specific to Scanlon's contractualism, due to the way probabilities of burdens are accounted when considered for rejection.

<sup>17</sup>This claim has been disputed by Alison Hills, who suggests that when obligations towards animals are taken into consideration the demands of contractualism become less demanding and more plausible than those issued by consequentialism (2010, p.241).

contractualism appears susceptible to the same criticism.

Let us consider the type of case which is often the focus of demandingness; obligations to the world's poor. As a test case, take Singer's Sacrifice principle,<sup>18</sup> which states:

“if it is in our power to prevent something very bad from happening, without thereby sacrificing anything of comparable moral importance, we ought, morally, to do it.”

(Singer, 1972, p.231)<sup>19</sup>

In order to determine whether or not this principle can be reasonably rejected, Scanlon states that we must consider various standpoints. Primarily, we focus on the “specific individuals affected by specific actions” (1998, p.202). With regards to this principle, this would presumably include the would-be victims of the very bad result and the would-be sacrificers. In addition, we must consider a broader picture, taking into account “the consequences of general performance or nonperformance of such actions and other implications (for both agents and others) of having agents be licensed and directed to think in the way that that principle requires” (1998, p.203). This requires us to think of how the general rule may affect others who could be affected by the principle outside of those immediately affected.

If *Sacrifice* is accepted, many affluent western people will be required to sacrifice goods to the extent that their sacrifice is not comparably morally important. This will presumably require *at least* the sacrifice of goods which have no (or negligible) effects on their wellbeing. Accepting that this will be considerably more sacrifice for many people than already takes place, the acceptance of this principle would result in many people in the developing world not suffering premature deaths and lower quality of life.

Because their interests would be harmed by the principle being imposed (even though it must be to a non-comparable degree), the affluent westerners

---

<sup>18</sup>Scanlon actually discusses this principle explicitly, calling it the “Rescue Principle” (1998, p.224). He claims that – once supplemented by a clause preventing iterations – it seems plausible that this could not be reasonably rejected.

<sup>19</sup>Obviously, what is demanded by this principle will depend entirely on what is regarded as morally significant, and if possession of even slightly more wealth is regarded as morally significant, it will only be negligibly demanding. In what follows I will assume, as Singer does, that this principle would require “a great change in our way of life” (1972, p.241) for typical western agents.

would have an objection to *Sacrifice*. Having an objection does not automatically provide a reasonable rejection though, as reasonable rejection is a comparative notion. If all other candidate principles gave others stronger objections, where those others must be understood as *individuals* (i.e. any would-be beneficiary of the donations required by the principle), then the westerners would not have a reasonable rejection.

As the types of objections to less demanding principles – those that do not require substantial donations – from those in the developing world seem to be far more burdensome than the objections that the affluent westerners have to the more demanding principles (like *Moderate Sacrifice*, or Miller’s *Sympathy*), contractualism seems to favour the demanding principles.

Reasonableness, in some of the ways described in section 6.1.5, does play a big role in determining what counts as a good objection. If we recall how reasonableness requires acknowledgement of the moral force of any individual agent’s goals or desires, this is clearly catered for in Scanlon’s account. It is the way the values of individuals are recognised by contractualism that leads to its denial of aggregation. Treating the interests of an individual as only a part of a group’s interests fails, in the eyes of the contractualist, to see each agent’s interest as important qua *their* interest.

One advantage that contractualism claims over utilitarianism (as well as its ease in recognition of the separateness of persons) is its ability to acknowledge reasons other than those pertaining to pleasure or pain (or welfare simpliciter). We might assume that acknowledging other reasons could provide grounds for rejecting other principles. As well as the burdens to welfare of donating resources to charity, the western agents find their property rights threatened and a burden of obligations supposedly incurred. However, the reasons held by each potential westerner for rejecting a demanding principle like *Sacrifice* would – in order to ground a reasonable rejection – need to be judged to ground a stronger complaint than that held by those suffering from easily curable illnesses in the developing world as a result of a less demanding principle. If, like Ashford, we accept that a “cost to an agent imposed by all but exceedingly demanding principles of aid will be outweighed by the cost to individuals” in dire straits (2003, p.288), Scanlon’s account still seems very demanding.

Despite accepting this feature of reasonableness into the account, it appears that Scanlon’s contractualism still leads to extremism. In the next

section, I attempt to demonstrate that an appreciation for the context-sensitivity of reasonableness can help to prevent this.

### 6.3.2 Problems in Scanlon's Contractualism

If a reasonableness-based account of obligation is able to contribute to a moral theory immune (or at least more resistant) to accusations of over-demandingness, then there must be an explanation for why the most prominent such account is unable to do so. After all, if I intend to suggest (as I do) that an account of moral obligation based on reasonableness can avoid the demandingness objection, why does Scanlon's – an obvious candidate as a sophisticated account which explicitly places reasonableness at its forefront – fail to do so?

I suggest that one reason this fails is due to the account of reasonableness that Scanlon provides.<sup>20</sup> Scanlon requires that all parties want to make an agreement, and that they hold a similar notion of what is reasonable. This raises a problem which is discussed by Margaret Moore, who notes:

“[Scanlon's] assumption that each person constrains her first-order desires according to what satisfies the ‘reasonably reject’ test is functionally equivalent to assuming a common standpoint for all moral subjects, which serves as the basis of agreement. The assumption is that each person reasons in the same way, and indeed that this reasoning is contractarian in form; thus, the contractarian argument makes sense because it mirrors the form of moral reasoning which occurs in each of us.”

(1996, p.176)

Moore contends that the presumptions that all parties reason similarly and *want to* reach agreements are unwarranted. Expanding and unpacking this critique further, we might consider two assumptions that Scanlon makes:

1. The assumption about what particular features make an complaint reasonable.

---

<sup>20</sup>This section draws partially upon criticisms of Scanlon (as well as Rawls and Barry) made by Margaret Moore (1996). While she made these criticisms several years before the publication of *What We Owe To Each Other*, they still apply. The same criticisms apply for Rawls's account, so I shall not argue separately for the claim that his account should also be rejected.

2. The assumption that all reasonable parties seek agreement with one another.

Let us begin with a cursory remark on the second of these. Moore seems to deem this a particular problem for Scanlon. She considers some of the biggest moral dilemmas facing contemporary society, like whether abortion is permissible or whether racially divided societies should permit group rights/protections or pursue an assimilationist agenda. Moore suggests that in such cases. . .

“it seems plausible to argue that [according to Scanlon’s account] one or both sides would be conceived of as unreasonable because unable to arrive at a compromise or policy which incorporates the desires of the other side”

(1996, pp.175-176)

When two parties are diametrically opposed with regards to an issue, it might be the case that they do not seek a justification that would appease the other side at all. One side may realise that any principle they would be willing to accept would upset the other side, and vice versa. Because “people have aims and interests and desires which they seek to advance in an agreement, and may reasonably disagree over what should be agreed to” (1996, p.176), but Scanlon requires for a person to be reasonable that they seek universal justifiability, Scanlon’s position must regard such people as unreasonable. As we all may have ‘red lines’ on some areas where we may not consider any compromise, Scanlon’s requirements appear problematic.

Rather than viewing parties like these as being unreasonable, we could suggest that they are being reasonable – they are attempting to justify principles in ways they think others should accept – but that there might not be a suitable principle about, say, abortion, such that “no one could reasonably reject” it. It might be noted that the implication of this, considering Scanlon’s account of wrongness, would be that morality (in Scanlon’s narrow sense) would simply be silent on issues of this sort. That response, however, seems unsatisfying, potentially excluding a vast amount of morality (and it does seem to be the ‘morality of obligation’, as Scanlon calls it).

Perhaps this is a resolvable issue. Scanlon could perhaps resort to higher-order principles about how people should act in these situations in order to

select the best first-order principles.<sup>21</sup> Regardless, I think the first type of problem may be more serious, so will set aside problems of the first type.

Returning to the first of the above reasons, Scanlon accepts that according to his account, “whether an action is right or wrong requires a substantive judgement on our part about whether certain objections to possible moral principles would be reasonable” (1998, p.194). In this respect, he seems to follow W. D. Ross, in thinking that even once all the empirical facts are settled, an element of judgment (perhaps something like intuition) is still required to find the correct verdict.<sup>22</sup> Because Scanlon requires there to be principles *no one* could reasonably reject, those reasons that can ground reasonable rejections must be accessible to all *and* everyone must be capable of making the same judgments.

Scanlon declines to specify particular features which may ground complaints, leaving these generally open, though acknowledges that well-being, will, in any good interpretation prove to be very important.

What other reasons may ground a rejection? Scanlon notes that unfairness is one such reason. A principle which arbitrarily favour one individual over another would be unreasonable. For instance, a principle which singles out a person by name would be unfair (1998, p.212). This seems in-keeping with the account of reasonableness given previously.

Scanlon also suggests that an agent’s previous contributions/efforts may make a difference, at least in providing a reason not to give up what one has achieved through those efforts.<sup>23</sup> We might note that this type of reason provides one way Scanlon could respond to the issue of iterative demands.

Another type of reason might be based on promises or agreements. If

---

<sup>21</sup>For example, Scanlon discusses the Principle of Established Practices (1998, p.339), which dictates that the status quo wins out when there is nothing else to decide between principles. This principle is discussed further in 6.3.3. Higher-order principles could provide a decision-procedure in cases like this, but challenges remain for justifying those principles, and ensuring that they give answers in each case (the Principle of Established Practices presumably is of no help when a new problem arises). While I do not consider this problem further, I am unconvinced that higher-order principles like this could give satisfactory answers here, and generally about the prospects of Scanlon being able to provide a satisfactory response.

<sup>22</sup>Ross elucidates this with the example of a poem. A poem can be beautiful or unpleasant on the basis of a variety of its features, but “our judgement as to the degree of beauty it possesses on the whole is never reached by logical reasoning from the apprehension of its particular beauties or particular defects” (2002, p.31).

<sup>23</sup>This is implied in the life jacket example (1998, p.196). Scanlon claims that of two agents deciding who should get a life jacket, it might make a difference to their objections that one *already* has the jacket, saying “perhaps he has looked hard to find it.”

one has agreed or promised to  $\phi$  in the past, this is also something that should bear on whether an agent has a reasonable rejection against one's  $\phi$ -ing or not  $\phi$ -ing (1998, p.200).<sup>24</sup>

All sorts of reasons could seemingly give an agent an objection, but there are restrictions. Scanlon claims that objections must be based on 'generic reasons'. These are reasons based on "commonly available information about what people have reason to want" (1998, p.204). The reasons must be based on what everyone has reason to want (in a general form), because Scanlon regards justifiability as crucial in accounting for the normative force of morality, and if the reasons were not available to some agent, they would not be able to provide a reason for them. Generic reasons for a person must be ...

"... based on the ways in which they would be affected by it: the ways their lives would be affected by living with the consequences of the actions it would permit and with the possibility that agents may perform such actions, since they would be permitted to do so."

(2014, p.96)

When discussing such reasons, Scanlon mentions that they obviously include "reasons to avoid bodily injury, to be able to rely on assurances... to have control over their own bodies" (1998, p.204). These seem largely uncontroversial. When we consider how generic reasons might feature in difficult moral dilemmas, however, it is unclear just what may be regarded as a generic reason.

Clearly Scanlon regards costs to welfare as extremely important. These certainly seem easily accommodated as 'generic reasons', as we take these to be based on "commonly available information about what people have reason to want", i.e. we know people have reasons to want to promote their welfare. We have already mentioned Scanlon's claim that an agent's having gained something through previous efforts may give them an objection to principles requiring that they relinquish it. Perhaps we can readily accept that by having committed significant energies to achieve or attain something,

---

<sup>24</sup>While Scanlon does think that promising can provide reasons of this sort, he does not think the practice of promise-making is special, rather seeing it as obligation-generating as a broader form of expectation-generating behaviours (1998, p.309).

that we have additional reason not to want to give it up (though I find even this extremely contentious).

Let us consider a difficult case. Two agents, A and B, have a moral dispute about the permissibility of some action. A has a generic reason which gives her an objection to principle *x*, and B has a generic reason which gives her an objection to principle *y*. For simplicity, assume these are the only two candidate principles. If A's objection is stronger, we should accept principle *y*. If B's is stronger, we should accept principle *x*.

Scanlon suggests that some cases like this might require a judgement from an agent. He gives an example of whether an agent would be required to sacrifice an arm to save the life of a stranger, and says he does not think "that any plausible theory could eliminate the need for judgements of this kind" (1998, p.225).

In the arm sacrifice case, Scanlon actually does not think the sacrifice of the arm is required, but offers no justification for this. The only consideration offered in favour of not requiring extreme sacrifice is the generic reason of intrusiveness. We can thereby reconstruct our dilemma. Agent A, whose life is at stake, has an objection to principle *x*, which states "If one can save the life of another by sacrificing one's arm, it is permissible not to do so."<sup>25</sup> A's objection is based on the reason that their life is at stake. Agent B, however, has an objection to the contrary principle "If one can save the life of another only by sacrificing one's arm, one *must* do so." B's objection is based on the loss of welfare from losing an arm *and* its being "intolerably intrusive" (1998, p.225). Perhaps B could also lodge an objection on the grounds of bodily integrity. Scanlon's contractualism then requires us to decide which of the objections is stronger. If, say, A's objection is stronger, principle *x* can be reasonably rejected. Scanlon takes it as given that this would entail that principle *y* could *not* be reasonably rejected, therefore it would be wrong to fail to sacrifice an arm in this type of case.

It is not fatal if Scanlon cannot provide a verdict in all of these cases. We could claim that in cases where it seems too close to call, or when the options are of similar magnitudes of importance,<sup>26</sup> no principle should be

---

<sup>25</sup>The principle would really need to be specified in various ways. We might want to specify things like the sacrifice being painless, that the agent's life would be worth living after losing the arm, that the other person is innocent and would probably live a long full life otherwise, etc.

<sup>26</sup>I add this caveat intending recognise cases with 'sweeteners', and not regard them

accepted and, as Scanlon suggests, an agent must simply make a judgement. It would be fatal, however, if it cannot provide a verdict in *any* case, or cases that we typically see as extremely obvious.

Let us now consider an apparently obvious case: gratuitous torture.<sup>27</sup> By this I mean torturing purely for the purpose of satisfying a sadistic urge of the would-be torturer. We might consider The Torture Principle:

**The Torture Principle** It is forbidden to torture any person merely for pleasure.

In order for it to be wrong to torture a person merely for pleasure, according to Scanlon, this principle needs to be such that no one could reasonably reject it (or that no one could reasonably refuse to accept it).<sup>28</sup> According to Scanlon's method, we must examine various standpoints, primarily those of the individuals affected by specific actions. The would-be victims of the torture are the main example here. Their welfare would presumably be incredibly enhanced by such a principle. Many people who might otherwise fear being tortured would also benefit from the security that acceptance of this principle would grant them. On the other hand, those would-be torturers will suffer from not being able to make themselves happier by torturing. The option to attempt to torture is also closed off, so even if someone did not really want to torture anyone, they might – and admittedly, I am considering people of particularly unusual psychological makeups here – feel constrained by the principle anyway.

Communities of people who love torturing and where this practice is accepted obviously are not commonplace, but we might imagine a civilization where a lot of people *really* enjoy occasionally torturing someone. Maybe the torturing takes a particular form which is deemed acceptable (but is still a horrific experience for anyone undergoing it).<sup>29</sup> People still dislike being tortured just like we do, but maybe they take extra relish in knowing that

---

as relevant. If, for example, the reasons in a situation seem to generate equally strong reasons in favour of saving one person rather than another, a discovery that one person is wearing a nicer shirt which would not be ruined if they are saved should not make the decision an obvious one. Many discussions of this sort are discussed in Caspar Hare's *The Limits of Kindness* (2013).

<sup>27</sup>Scanlon does include torture as something there is “no doubt” is wrong, so presumably he should be able to account for this (1998, p.357).

<sup>28</sup>Scanlon implicitly accepts that these are equivalent (1998, p.195).

<sup>29</sup>Perhaps we could imagine some strange regulations about the torturing, restricting its time or type, like in the film, *The Purge*, for example.

being tortured is horrific and everyone there (for simplicity) thinks it should be permissible because the prospect of getting to torture someone is just so appealing. Perhaps, we can even imagine some people in that community who do not plan on ever torturing anyone, upon hearing a proposal that a principle banning torturing was being considered, seeing this as overstepping upon individuals' freedoms. If they judge certain positive freedoms, like the freedom to occasionally torture someone in a particular way simply for the sadistic pleasure, to outweigh the cost of being (or potentially being) tortured, it seems that people in this community would want to reject the Torture Principle.

Given that the types of justification for rejecting the principle are similar in this case and the arm case (welfare and intrusion on freedom),<sup>30</sup> it might seem difficult to come to the apparently obvious conclusion prohibiting torture in one case, and not require sacrifice in the arm case (as Scanlon thinks is the case).

Many people will deny that this type of story I have given here is plausible, or choose to restrict our moral principles to circumstances that are sufficiently like the actual world, where torture-sympathisers are clearly not widespread. I picked an extreme type of example deliberately. For those who accept that this demonstrates my claim, that Scanlon's account cannot even tell us that gratuitous torture is wrong, my work is done. If his account fails to demonstrate this, it is unable to provide us with even the most uncontroversial moral principles, and must be deemed a failure.

For those who deny that the type of society I have described is possible, or think it is too removed from typical human relations to be useful, I offer a the following challenge. Consider the wide range of practices most of us judge to be abhorrent that have been deemed permissible throughout human history, such as slavery, subjugation of women and many cruel practices towards non-human animals. Some of these views plausibly were deemed permissible because of some factual mistakes, like the views that animals do not feel pain (or feel it differently in some important respect). Alternatively, with regards to some of these cases, we might think that the practitioners

---

<sup>30</sup>We could claim that the loss of welfare through losing limbs/life is different in kind to – and providing different reasons to – the loss of welfare through being tortured or not being able to partake in torture. While this might be the case, I do not think it detracts from the argument here, as the conclusions Scanlon wants to reach still require that we reason the same way.

did not *really* think what they were doing was permissible. In some of these cases, it seems that explanations of this sort are appropriate, but this does not always seem to do the work. Consider some instances of particularly unpleasant battery farming. It seems very plausible that some of the people who understand what these practices entail think that this confers an obligation to not be complicit in this process (not consuming those goods and possibly some activism to prevent), whereas others – seemingly given the very same facts – do not accept such an obligation. If we accept that people have the same information, and that they genuinely do have different permissibility-judgments, why would we not accept that this could be the case with principles like The Torture Principle?<sup>31</sup>

It seems problematic for Scanlon if even one person, on the basis only of generic reasons, judges that this principle should be rejected. There seem to be three responses which could allow Scanlon's view to give the verdict he wants. Firstly, the would-be torturer might have some reasons which are not generic reasons. Secondly, we could suggest that there are some additional generic reasons in any instance like this, which allow us to reject one principle. Finally, Scanlon might suggest that in each of these there is some way we might determine which principle should be rejected, and that it would be impossible to reasonably accept obviously mistaken principles.

The first two of these options do not seem viable. The reasons in play seem like those based on “commonly available information about what people have reason to want”; pursuit of welfare and freedom from excessive intrusion from morality. We might well imagine that there could be additional reasons in cases like those mentioned, but those would be artefacts of the examples, and could presumably be eliminated with more detailed cases. The best option for Scanlon thus seems to claim that it *should* be discoverable, simply from considerations of the principles, and the standpoints of those who could be affected by them, which principles must be accepted.

To achieve this, it seems that Scanlon requires “a common moral standpoint”, as Moore puts it (1996, p.176). The assumption is that each person

---

<sup>31</sup>This response is admittedly brief, and I do not expect that this will be convincing to everyone. However, at very least I think it offers a challenge. It does *seem like* people with the same information do arrive at different verdicts in this type of case. My opponent must either demonstrate that people who arrive at different verdicts either deny that these cases are possible (that agents must have different information, be engaging in faulty reasoning, or being disingenuous about their verdicts in these cases) or provide some explanation for why these cases are different to The Torture Principle.

reasons in the same way, insofar as they judge obligations to arise from reasons, is what I dispute. Fiona Woollard labels this – a “given way of assigning force to reasons in deliberating” – the *Reasons Weighting* (2016a). If what obligations are reasonable considering certain reasons varies, between communities for example, there will be no such common moral standpoint; there will not be a universal Reasons Weighting.

Scanlon argues that his account is no more problematic than one which involves only one type of reason (e.g. pleasure in hedonic consequentialism), extolling the benefit that he need not reduce a wide variety of moral explanations to one idea.<sup>32</sup> Even if he is correct – that weighing reasons of, say, welfare against freedom, provides no additional problems – this type of problem can persist even if we held a monistic theory. It is an additional question, after all, if one accepts that only welfare can provide moral reasons, when those reasons can incur moral obligations upon individuals. Maximising act utilitarians give one answer to this, that an act’s causing the most happiness makes it obligatory, making a theory of obligation easily derived from their axiology (a route I have suggested, e.g. in section 5.2.4, that we should reject). Satisficing utilitarians, however, give a different answer setting a standard one has an obligation to meet, and scalar utilitarians give yet another answer.

What does complicate matters in an account, like Scanlon’s, which embraces a pluralism of types of reasons which may contribute to the acceptability of a principle, is that before even arriving at the question of whether an obligation is formed, some weighing of these reasons against each other needs to take place before we can even know which actions are best. This is a problem faced by any pluralistic account, but even if it is resolved – and this would also seem to rely upon a common Reasons Weighting – the question of when obligations result remains.<sup>33</sup>

---

<sup>32</sup>Notably, Singer and de Lazari-Radek do think this is a benefit for monistic theories. They suggest that an account with simply one type of good, like a Sidgwickian hedonistic utilitarianism, does have an advantage – at least in principle – over pluralistic accounts like Ross’s because it can “give a definite answer to the question of what I ought to do, whereas on Ross’s view, when duties conflict, one can only intuit, in each set of circumstances, which duty or duties carry more weight” (2016, p.206).

<sup>33</sup>Some accounts can avoid these problems. If we hold that some reasons are conclusive, then they will override other reasons in the first stage of this and guarantee some obligation. This would be the Kantian position. Kant holds that there can be no conflicts of duties, because in situations where there are multiple grounds of obligation, the “stronger ground of obligation prevails”. The weaker ground therefore can not form a duty, and is

If there is no common Reasons Weighting – no objectively privileged weighting and obligation-generating method<sup>34</sup> – there may be no principles such that *no one* could reasonably reject them.<sup>35</sup>

It should be noted, however, that a relativised notion of reasonable rejection – which could be shared among agents who accept similar interests as grounding rejections and similar force for these reasons – would still be a live option. This shall be returned to in section 6.4.

### 6.3.3 Why this matters for Demandingness

As has previously been discussed, because of problems of iteration (section 2.2.1), accepting a principle like Singer’s *Sacrifice* makes a theory extremely demanding. Scanlon suggests that a similar principle to *Sacrifice* would be acceptable under his contractualism. However, he adds a caveat that previous contributions are a factor that should be taken into account, so that such a principle could not require “unlimited sacrifice” (1998, p.224). Ashford argues that through the same reasoning Scanlon makes use of justifying *Sacrifice*, we could also be committed to accepting what she calls ‘The Stringent Principle’:

**The Stringent Principle** If we can prevent something very bad from happening to someone by making a great sacrifice (e.g. giving most of our income to aid agencies and spending a lot of our spare time on campaigning and fund-raising), it would be wrong not to do so.

(Ashford and Mulgan, 2012)<sup>36</sup>

This would make Scanlon’s account extremely demanding without even requiring iterations. Ashford suggests that the contractualist should be committed to this stronger principle because the “sacrifice imposed on an agent thus excluded from the moral calculus (1996, VI:224). Of course, this type of picture needs a way of determining the ‘stronger ground’, and similar difficulties are likely to ensue in that project.

<sup>34</sup>I use the term “method” to be as theory-neutral as possible here. For instance, particularists like Dancy (2004) or anti-theorists like Williams (1973), would reject the notion that any sort of algorithmic system would be possible, but would accept that we have some method of evaluating

<sup>35</sup>Thomas Nagel suggests that this might be the case, and that in situations where there are no solutions that “no one could reasonably reject, neither party... can be reproached for trying to impose a solution acceptable to him, but unacceptable to his opponent” (1991, p.173).

<sup>36</sup>Originally argued for in Ashford (2003, p.287).

by a principle of aid to those in need will have to be extreme before it balances the cost faced by individuals in dire straits who are not helped and otherwise would have been" (2003, p.288).

At this point however, it seems that the strength of the complaint – generated by the generic reason based on welfare – is equated with the severity of the costs. It is natural for us to be lured into this manner of evaluation because accounting for costs to welfare is such an important part of our moral deliberations.

Herein lies how contractualism often appears extremely demanding. Absent some weighting criteria for how various reasons should come together to form our obligations, we may simply consider what is most important. I want to suggest that this drives us in a different direction than when we consider what obligations we think are reasonable. Some principles/practices are internalised within a moral community, and these play a role in what is deemed obligatory. If we want an account that corresponds – even loosely – to our intuitions, we need to appreciate how the moral significance which we might place on them relates to these principles/practices.

Thus far, I have noted a challenge to Scanlon's account, namely that of justifying the claim that there exists some objectively privileged Reasons Weighting which grounds our obligations. An argument can be made that the existence of such an objectively privileged weighting (at least alongside some other assumptions Scanlon makes) is incompatible with some of our moral intuitions. If this is the case, we are confronted with the choice of either rejecting part of Scanlon's framework (which I suggest) or accepting that our intuitions are extremely mistaken.<sup>37</sup>

I contend that Scanlon's account leaves us unable to explain some of the moral phenomenology previously discussed. Let us suppose that Scanlon's account could give us a good Reasons Weighting that allows for an evaluation of a plurality of morally relevant factors. Even then, could an account like Scanlon's satisfactorily explain the different obligation phenomenology across moral communities?

If we consider the intuitions in Le Chambon compared with our own, they judge that some behaviour (harbouring refugees) is obligatory, whereas we judge it as optional (supererogatory). It could be suggested that one of the groups is simply mistaken about what is obligatory. I take it that this is an

---

<sup>37</sup>The latter of these options is considered in section 6.4.5.

undesirable conclusion. If we accept that we are mistaken about this, we are likely to deem the demandingness objection unproblematic.<sup>38</sup> It seems chauvinistic to suggest that we are correct and villagers of Le Chambon are mistaken, particularly since the types of evidence available to us seem to be the same as the evidence which was available to them (considerations of the values in question and their moral intuitions).

It seems plausible that there are similar (though perhaps often less drastic) differences between a great many communities. We might again consider the example of a community that regards recycling plastics as compulsory. Another might accept all the good-making features of doing so, and have similar facilities, but not regard this as obligatory. Considering that changes in our judgements of the moral landscape – what we judge to be acceptable/unacceptable behaviour – are fairly frequent,<sup>39</sup> it also seems difficult to claim that our obligation-intuitions are the right ones.<sup>40</sup> At first glance, it might appear that Scanlon is unable to explain this other than by claiming that one of these communities is mistaken.

However, there are a few potential options for Scanlon. Firstly, he does offer an explanation for how conventions can affect what is required in a situation. One such explanation comes in the form of the “Principle of Established Practices” (1998, p.339). This holds that in certain situations, where there would be no way to decide between numerous otherwise non-rejectable principles, if one is generally accepted within a community (which need not be unanimous), this gives an additional reason not to reject that principle, i.e. it would be wrong to violate the generally accepted principle in favour of some other.

While this higher-order principle would allow for acceptance of different principles in different communities, it does not seem to help in situations like Le Chambonnais case. This is because, presumably, the refugees facing very likely death have a stronger objection to the principles that permit not

---

<sup>38</sup>Chapter 2 discusses why I reject this option.

<sup>39</sup>Lindström et al. provide many examples of changes in permissibility verdicts in a range of issues, from tax-dodging to free-riding on public transport (2017).

<sup>40</sup>We might be tempted to think that something like scientific progress has occurred in the moral domain, and that we could reject former moral verdicts the same way that we would now reject bad scientific theories. This option, however, is problematised by the feature that in some of the moral cases it is agreed exactly what the evidence is, and we may have no way of demonstrating why our Reasons Weighting is superior – we might think that in instances like the Chambonnais case, their weightings seem preferable in some sense.

coming to their aid than do the locals, so the higher-order principle would not come into play.

Alternatively, we could find another way that Scanlon might respond to the (apparently) contrasting verdicts about obligations here in his contractualist formulation of wrongness:

“It holds that an act is wrong if its performance under the circumstances would be disallowed by any set of principles for the general regulation of behaviour that no one could reasonably reject as a basis for informed, unforced general agreement.”

(1998, p.153)

If we take the restriction of *under the circumstances* very broadly that could make a difference here. He discusses this option for how the same types of act might be permissible in one community, but forbidden in others. He does this by reference to the reasons people have, noting that “what people have reasons to want depends on the conditions in which they are placed, and among these conditions are facts about what most people around them want, believe and expect” (1998, p.341). This certainly coheres with the considerations raised in chapter 3 about how these could be morally relevant.

The type of variation in what is required by principles which Scanlon permits is accordance with what he calls *parametric universalism*, which is compatible with his form of contractualism. Parametric universalism consists of two claims:

1. There is *one* single set of standards with moral authority.
2. These standards allow for variation of what conduct is right depending on circumstance.

(Scanlon, 1998, p.329)

While parametric universalism permits variation in right conduct, the differences in obligations must result from differences in reasons. He elucidates how different circumstances could lead to different actions being obligatory in different societies with an example about privacy. Scanlon notes that in societies with different forms of commerce, or where there are different ideas about dignity accepted, people will have different reasons available

to them. In our society where it is deemed embarrassing to be watched doing certain activities and where it is financially disadvantageous for people to know various personal financial details, we have reason to desire “forms of protection of the sort that rules of privacy provide,” but societies lacking these features do not (1998, p.341).

We might therefore think that varyingly demanding moral prescriptions could be accounted for similarly, by people in different communities having different reasons. For example, he could say that in Le Chambon, *the circumstance that* people typically think ‘one should help in these situations’ contributes to the evaluation, giving the members additional reasons to assist, making it obligatory.

However, this is not available to Scanlon for at least two reasons. First, he sees it as a problem of relativism that should be avoided that obligations and moral entitlements might not properly match up (1998, p.332). If he allowed for this interpretation of his account, it seems like it would be possible to have, for example, refugees entitled to help – that is, able to legitimately claim to be owed assistance – in Le Chambon and in a neighbouring village, but the locals only under an obligation to help in one of these villages. Scanlon regards it as a problem if we can accept that an individual legitimate claim on another to be owed something, but the other has no such obligation.

Secondly, Scanlon requires that the agents have different reasons in order for the different obligations to obtain. While he accepts that identifying with different ‘ways of life’ could provide these reasons, Scanlon explicitly rejects the notion that such a reason could be of the form “I have to do this because the custom in my community requires aiding refugees in this case, so I am obligated to help here”. Scanlon regards this as failing to capture the moral motivations, noting that when people give justifications derived from on customs or tradition “reasons of the kind in question lose their force” (1998, p.336). The reasons would have to come from a special meaning or value, not merely a recognition of custom, to be regarded as providing the appropriate moral motivation. However, it seems implausible to suggest that, were we placed in a situation like those in Le Chambon, we would have significantly different reasons available to us. The reasons to help (to prevent the suffering of the unfairly persecuted) and not to help (not wanting to accept the risk to one’s personal safety) seem the same in

these cases.

If we accept a pluralism of systems of how reasons (even accepting Scanlon's generic reasons) could be weighted in the generation of obligations, we could escape the demanding conclusion. Providing a schema for how such a system might operate and attempting to defend something of that form is the aim of the remainder of this chapter.

## 6.4 Relative Reasonableness

The main problem in Scanlon's contractualism, in terms of demandingness, was that it asks us to act in accordance with principles no one could reasonably reject, and assumes we can discover those principles based on the strength of generic reasons. I claimed that when evaluating what obligations are prescribed by a theory, in lieu of some 'real' objective strength of these reasons, we are likely to select principles likely to avoid certain features which strike us as particularly bad. Scanlon does place some restrictions on these – for instance, impersonal reasons cannot count, and the reasons must be reasons *for* an individual – but we are still making evaluations based on judgements of value, rather than about justifications.

That we fall into a trap of considering welfare is indicated by the way people (including Scanlon) talk about these cases – the discussion often looks more about *cost*. For example, when considering whether a potentially intrusive principle could be rejected, Scanlon says we ...

“... do not just compare the costs ... of abiding by it, or not doing so, on a specific occasion. We have to consider also the general costs (and benefits) of its acceptance.”

(1998, p.225)

While some generic reasons are easily understood as costs – like a loss of money or bodily integrity – not all are. Just as difficulty is different to cost, many of Scanlon's reasons can not easily be regarded as costs. Having already contributed a significant amount, for instance, seems difficult to understand as making giving more now costly. Even fairly clear reasons, like having made a promise, seem strange to understand in terms of *cost*. If we do try to understand them as costs, they seem less obviously burdensome

than the truly obvious costs – like loss of life or suffering through ill-health. Because of this, we might fail to do justice to our actual Reasons Weightings when we think this way. Ashford also talks largely in terms of costs when she tries to replicate Scanlon’s reasoning to reach her Stringent Principle. For example, she claims that “the cost to an agent imposed by all but exceedingly demanding principles of aid will be outweighed by the cost to individuals in need imposed by less demanding principles of aid” (2003, p.288). In addition, thinking in terms of cost-benefit analysis and accepting ‘the best’ seems in danger of leading us to the type of presumption of maximisation that was identified as problematic in the previous chapter.

I suggest that because of a lack of an obvious method, we consider *obvious* costs, and might be misled into overly demanding accounts. Of course, Scanlon could respond that there is a *true* Reasons Weighting, and that when this is adopted, and we think about cases properly, we will get the right verdicts. However, if there is an objectively privileged Reasons Weighting, we are left with a puzzle. Consider a case where we hold that some behaviour (in some circumstance) has a certain obligation-status. We also know that other people, from a different moral community, believe the same behaviour has a different moral status – we might consider saving refugees in Le Chambon, which we deem supererogatory, but the residents of the village deemed obligatory. It also seems that we do not have any better access to the truth of the matter than our rival moral community does. The problem is how to reconcile our belief about the moral status of the behaviour with our apparently conflicting belief that the other community is just as likely to be right as we are.<sup>41</sup>

One solution is to accept that obligations are generated by Reasons Weightings, but that no Reasons Weighting is *correct*. This is to accept some relativism into our verdicts of moral obligation. In the remainder of this chapter, I attempt to flesh out further how such an account might work. This is still a superficial sketch, but I hope to demonstrate some advantages that such an account might possess.

Let us reconsider the example of Le Chambonnais from chapter 3. If we

---

<sup>41</sup>Of course, one could simply deny that the townspeople were under an obligation in this case, and suggest either that they were mistaken (that the acts were actually supererogatory) or that they were being disingenuous in their reports of how they viewed their acts. I do not consider this in this section, but this option does seem hard to maintain.

make the move – which I previously endorsed – of accepting that the actions and expectations of the moral community served to modify the obligations of the townspeople, how might we explain this?

One explanation offered in that chapter was that the compliance and expectation of the community changed how difficult those actions were. Once everyone else acted in a certain way, mental barriers against the noble behaviour were removed and/or it became more difficult to not act nobly when this was the more readily accepted behaviour.

In the chapter 5, this explanation was rejected, as difficulty was argued to only play a role in providing an excuse condition. We could still suggest that what occurred as more and more of the village accepted the burdens of aiding fleeing Jews, was that it would have been less difficult, so failing to do so would have been less excusable for them than for us. This still fails to capture some important phenomenology though, because we do not even judge that failing to help in this case *needs* an excuse. Unlike the villagers of Le Chambon, we simply do not see this as obligatory. Also, from our perspective, when told of the villagers who did not risk their livelihoods, we are likely to judge them no more harshly than the villagers of any other village.

Instead, I suggest an explanation incorporating reasonableness. We have noted that the reasonableness of certain other-regarding behaviours depends on what is “ordinary” or “customary”. This will allow us to accept that social factors like expectation and compliance affect moral obligations. However, whereas in chapter 5, we attempted to identify the way these factors operate via difficulty, a more successful approach may be through our notion of reasonableness, which seems partly constituted by what is ordinary or customary (in some sense).

If, in Le Chambon, it was ordinary or typical (or even perceived as such by a large enough number of villagers) to risk one’s life in situations like those that obtained, could that make it reasonable to require that behaviour? Recall the tentative definitions offered earlier:

**$\phi$ -ing is impermissible if and only if** it is reasonable to require an agent not to  $\phi$

where  $\phi$ -ing is an object of moral judgments.

Rather than making use of generic reasons to explain what it is rea-

sonable to require agents to do (or not do), we can thoroughly exploit the notion of reasonableness depicted throughout this chapter.<sup>42</sup> What it is reasonable to require people to do will depend upon many contextual factors, but these factors will be those accepted by the moral community. For example, if one's moral community judges that the dire situations of refugees impute moral obligations upon those able to assist, it might be obligatory in that community to offer assistance (as might explain the phenomenology in *Le Chambon*). Similarly, if in a given moral community, having rendered assistance many times already is accepted as alleviating some obligations to help in future, iterative demands might not make result in the extremely demanding burdens suggested by Singer et al.

This account does not need to permit a total moral 'free-for-all', wherein any rules, no matter how unfair or absurd can be accepted. There will be constraints concerning what principles can be adopted, based upon what is accepted as grounding obligations. If one accepts something like a nearby rescue principle<sup>43</sup> then they must accept the verdicts this entails regardless of whether they are the one in need of rescue or the one in a position to do the rescuing. Reasonableness seems to require that all principles governing obligations can be justified agent-neutrally.

There are many questions to be answered with an account like this. Many concern the nature of the moral communities. For example, one might ask:

- How big are moral communities? (Can there be a moral community of one?)
- How are moral communities formed?
- Is a moral agent in *one* moral community, or can they be within several?
- How do people from different moral communities engage in moral discourse?

---

<sup>42</sup>The use of generic reasons could be preserved if this were deemed desirable, but where the differences in verdicts is accounted for entirely by the Reasons Weightings. However, I will not take this option here, so to allow for contextual features to play a role in what is deemed to have value.

<sup>43</sup>For example, Richard Miller offers a plausible nearby rescue principle he thinks everyone must accept, that an agent "has a duty to rescue someone encountered closeby who is in imminent peril of severe harm and whom one can help to rescue with means at hand, if the sacrifice of rescue does not itself involve a grave risk of harm of similar seriousness or of serious physical harm, and does not involve wrongdoing" (2004, p.378).

- How complete will the theory of requirements be within a moral community?

Each of these questions has many plausible answers, which merit independent exploration, and much more discussion than I can provide here. To provide a thoroughly unsatisfying snapshot of how I would envisage the moral communities, here are some brief answers: I think the moral communities may be of *any* size (even one), so long as blaming and praising practices can take place. I also suppose it is via these practices that moral communities emerge. The notion of moral sub-communities seems entirely plausible. We might consider certain religious sects or groups of vegan friends within a wider community who would blame and praise each other differently (either more/less stringently, or for different behaviours). How moral communities might interact is a problem, but typically these communities will have significant overlaps in what they expect/require. Communities will typically have an ‘open-textured’ moral theory. By this, I mean that there will be possible cases outside the standard domains where the theory does not give a determinate answer;<sup>44</sup> their theories will include requirements/prohibitions mainly for common occurrences, and when new situations arise, perhaps with scientific developments which create new moral dilemmas, moral communities must adapt accordingly.

The relativism entailed by this proposal is often deemed undesirable. I argue why I think some of these problems are less of a problem in this picture in section 6.4.3. First, we can examine how this type of account can resolve the issues relating to demandingness that we have come across with regards to alternative theories.

#### 6.4.1 Conclusions so far

In this chapter so far, I have suggested that incorporating the notion of reasonableness into our accounts of obligation-generation might be helpful in giving an accurate – and not overly demanding – picture of what we take our obligations to be. Given that Scanlon provides an account that relies heavily on reasonableness, I explored his contractualism in particular.

---

<sup>44</sup>The term “open-textured” is attributed to Waismann (1968), though he does not give an explicit definition. Tanswell offers the following definition, which I will accept: “A concept or term displays open texture iff there are cases for which a competent, rational agent may acceptably assert either that the concept applies or that it disapplies” (2017).

I claimed that the assumption of an objective Reasons Weighting – which his account requires – is unjustified. I speculated that because ‘strength of reasons’ is a fairly opaque notion, and because we have no obvious method for comparing complaints, we are often compelled to think in more obvious axiological terms, which leads us to consider costs to welfare, rather than many features we do ordinarily take to be morally relevant. I suggested that this is implied by the way that talk about ‘strength of reasons’ often makes way for talk of costs, which does easily accommodate taking a variety of reasons as important.

Acceptance of an objective Reasons Weighting also leads us to a difficulty in explaining cases where different moral communities have different obligation-verdicts, despite seemingly having the same evidence. There are certainly ways that Scanlon’s account could be modified in order to avoid this problem, but I have instead suggested that we consider a relativistic Reasons Weighting; that different moral communities view the obligation-generating powers of certain reasons differently. What it is reasonable to demand in a community will depend, at least in part, upon issues of expectation and compliance. This is motivated by cases like Le Chambon and the recycling case.

We may then say that some act,  $\phi$  is impermissible (morally wrong) within a moral community if, within that community it is reasonable to require them not to  $\phi$ . What is impermissible in one community may be permissible (or potentially even required) in another, depending on their Reasons Weighting (which in turn is dependent on factors relating to compliance and expectation). This accepts an element of relativism, which I address in section 6.4.3. Prior to that, I attempt to demonstrate how the picture I have provided is able to respond to some of the challenges of demandingness.

## 6.4.2 Other Issues of Demandingness

Having given a schematic for an account of this sort in order to avoid over-demandingness, an obvious question is whether it actually does succeed in doing so. Because what a given moral community deems reasonable dictates what it judges to be required of a moral agent, there is a neat tie-up ensuring that agents are not systematically under what they view as unreasonable obligations (they in fact, definitionally, cannot be under unreasonable obli-

gations).

### ***Underdemandingness?***

This answer might seem a little unsatisfying. An uncharitable reading of this might seem to suggest that whatever you believe about your obligations, you are correct, and that this might actually make the theory *underdemanding!* In response to that it should be reiterated that this is not an ‘anything goes’ account. For instance, it does not entail that people are always correct about their obligations, because the picture is compatible with moral mistakes. These could come about in a variety of ways. People may make mistakes about facts. They might believe, for instance, that a member of another race does not have feelings of the same kind as them. From this, they may believe that the desires of those people count for less (in terms of generating obligations) than their own do. It might be that they are culpable for this mistake, even by their own lights, as people sometimes fail to scrutinise their beliefs sufficiently.

Another way people may make moral mistakes is via what Kant called *rationalising* [vernünfteln]. Sticker and van Ackeren claim that Kant saw rationalising as . . .

“... an abuse of rational capacities, which finds or invents apparent justifications or apparent excuses for violations of the moral law.”

(2015, p.86)

We could imagine agents rationalising by overestimating how important something is to them, fabricating reasons that or underestimating how much others value something.<sup>45</sup> This seems conceivable as a deliberate (or even semi-deliberate) way that people might come to moral mistakes. Something similar might happen entirely accidentally though. Because we have much greater access to our own needs and desires, and each of us is particularly focused on our own lives, we may (blamelessly or not) fail to appreciate the needs or wants of others, and how we should (given our Reasons Weighting)

---

<sup>45</sup>The ways Kant discussed the possibility of rationalising seemed to focus more explicitly on the nature of morality. Sticker and van Ackeren included questions whether morality is overriding, doubting the purity of morality or challenging the normative force of the moral law (2015, pp.85-86).

see this as affecting our obligations. This may also lead us to making moral mistakes about what is required of us.

Another way a moral mistake may arise is if the standards one takes to determine what is reasonable are inconsistent. If an agent held the belief that lying was wrong (that the reasons against lying were always *concluding* reasons), but also that doing anything to bring about the deaths of many people was wrong (also as a conclusive reason), they could fall into an inconsistency if placed in a situation where the only way to avoid causing deaths to many people is by lying. I provide an example like this to demonstrate that such inconsistencies might be mere ‘practical inconsistencies’ – most people will never encounter a situation like this, so might operate with an inconsistent view of obligation forever. I suggest that holding these views without noticing that they are inconsistent, even though the inconsistency is usually harmless, is still a type of moral mistake.

Of course, it is *possible*, within the framework I have provided, to have a moral community for which no reasons are deemed to generate any obligations. Such a community would regard no actions (or omissions) as morally wrong. This definitely would be an undemanding. However, while *conceptually possible*, such a community would be extremely different to anything we are familiar with. Due to the conceptual link between wrongness and blameworthiness, such a society could not regard any acts as blameworthy.<sup>46</sup> They could still make assessments, e.g. that some action is better or worse than another, but moral discourse of this sort would be, by our standards, extremely impoverished.

### **Cullity’s Iteration Problem**

In earlier chapters we discussed a problem posed by Cullity. I suggested that this poses a serious problem to certain theories, including those incorporating Scheffler’s agent centered prerogatives. The problem facing that sort of theory was that no matter what weighting I give my own interests, either the weighting will be wildly implausible – I will practically regard my interests as *far* more important than the interests of others – or I will still be left with a very demanding conclusion.

Our typical pre-theoretic intuition suggests that people have to con-

---

<sup>46</sup>At least, not if we maintain the independently plausible thought that an act can only be blameworthy if it is morally wrong, discussed in section 5.1.

tribute something towards certain needy charities (if they can afford to), but that once they have given a certain – and not extremely demanding – amount, they can stop (at least for a time), and that doing more than this is supererogatory. Scanlon’s account seemed able to deliver this verdict if some weight is given to the reason of having previously contributed (so long as an appropriate Reasons Weighting is given).

Some moral communities might accept that an agent who has already contributed significantly to some project (in time, effort or resources) is under no obligation to do any more. This certainly seems true of the moral community we are in. There are definitely situations where we accept that people can sit back having “done their bit”. This of course does not mean that obligations would be alleviated in every domain. While we might say of a soldier who has spent several years at the front line, that they are under no obligation to go back, we might always expect people to do their recycling, regardless of how much recycling they have done in the past. Similarly, if someone has just given a large portion of their money to an effective charity, we may judge that they need not contribute more (for at least some time), but if an agent is walking home and finds a stranger in dire need of medical attention, we judge that they have some obligation to help, even if it so happens that they just saved a similarly maligned person just the day before. What counts will be what reasons are in play and the extent to which one’s moral community sees those reasons as generating obligations. Other moral communities, of course, might have different requirements, but these do not govern our conceptions of obligation, or the blaming/praising practices that go along with them.

### **Singer’s challenge?**

Singer suggests that his principle, which he notes appears “uncontroversial” (1972, p.231), demands significantly more of people than ordinarily expected. We are then left with a dilemma of either rejecting Singer’s principle, or accepting that morality is much more demanding than we intuitively expect. In section 2.2.2, I suggested several reasons why we should reject Singer’s principle. I claimed that it was poorly motivated and that it leads to counterintuitive verdicts in a wide variety of cases, in addition to the case of our obligations to the global poor.

The blueprints provided above provide resources for not only why Singer’s

principle should be rejected (or at least, need not be obviously accepted), but also why any ad hoc revisions to the principle will also fail. If what is obligatory within a moral community depends in part upon what is typical, customary or accepted in that community, then it is no surprise that extremely unusual acts of beneficence (like those Singer judges as required) are not obligatory. If an agent's community do not view the existence of suffering that can easily be prevented as generating an obligation to help – if the Reasons Weighting of this type of reason is insufficient – that agent will not judge that they are required to aid. It is of course possible that some moral community could accept some very demanding moral principles (and by their standards, what we regard as ordinary behaviour might well involve flouting obligations), but we are not in a moral community like that.

Singer could reply that there is an inconsistency in our attitudes between cases like saving drowning children in ponds and those regarding our ability to help the global poor. The purported inconsistency here would be an inconsistency of responses to a certain type of value. However, there are other factors we may regard as important in generating obligations, which an analysis of that sort will not capture. Our community may deem that some sort of connection (social, psychological or causal) that holds between us and the drowning child is of the sort that has additional obligation-generating power. Alternatively,<sup>47</sup> our community might value our ability to plan our lives, perhaps seeing this as integral to our autonomy, and thus judge that events we can typically expect to occur either never or rarely have a greater pull on us.

The drowning child and the child dying of the preventable illness may be accepted to be equally *bad* in terms of sheer welfare,<sup>48</sup> but the norms and conventions we identify with in our moral community mean that we find allowing the former forbidden, but the latter permissible. Because of the conventions in our moral community, obligations are generated in typical nearby rescue cases – it is reasonable to require agents to act – but not (normally) in distant poverty cases.

---

<sup>47</sup>As was suggested in the discussion of Unger in section 2.3.2

<sup>48</sup>Again, I do not want to commit to any axiology here. A simple hedonistic axiology will only count the welfare differences, so would determine simply based on the pleasures/pains in question. If, we have even a slightly more complicated axiology, for example seeing pleasure *and* virtues as intrinsically valuable, as discussed by McElwee (2015), would complicate matters. Rival axiologies will provide different verdicts to what acts (or states of affairs) are morally better/worse.

### 6.4.3 Relativistic Problems

Moral relativism is often seen as objectionable, abhorrent or even scary. Scanlon describes three worries that are taken by its opponents to be reasons that we should fear or resist moral relativism. These are that moral relativism:

1. creates moral monsters
2. delegitimises moral condemnations
3. undermines moral authority

(Scanlon 1998, pp.330-333)

Scanlon suggests that the first of these should not really worry us. The worry that moral relativism creates moral monsters follows from the thought that if some of (what we judge to be) the most obvious moral verdicts fail to apply in all cases, then everything must be permissible (to at least some people). The people within those context where, for example, the prohibition against torture does not apply, would have *no reason* not to torture. If moral relativism were to spread as a doctrine around the world, one might worry that more ‘moral monsters’ would be created, judging that moral requirements – not being universal – do not apply to them. Scanlon seems unconcerned by this worry, claiming that historically the “worst mass murderers have not been relativists, and many relativists accept, perhaps for varying reasons, the basic contents of ordinary morality” (1998, p.331). Of course, this is an empirical matter and might be speculated about, but regardless, this worry seems more to do with the desirability of sharing moral relativism with the masses, rather than the merits of the account.

The worry that moral relativism makes moral condemnations less legitimate becomes apparent when we consider agents who have committed (what we take to be) horrible acts at various points in history or other geographical locations (some time/place the acts would not be forbidden). If moral relativism is true, it seem that those agents had no reason not to act as they did, and if that is the case, it seems inappropriate to criticise them for their actions. More worrying, Scanlon suggests, is that if we are unable to criticise an agent for violating some supposed right of another, we are not able to defend the claim that the other was entitled *not* to be treated in

that way. A form of relativism must be able to either [1] explain why losing (or, more accurately, that we do not possess) the ability to legitimately condemn such actions is not theoretically problematic, or [2] demonstrate that legitimate condemnation is possible under the relativistic framework. I opt for the second of these options in the next subsection.

The third concern about moral relativism relates to the authority of morality. Under the framework I have described, it is possible that a practice which is forbidden in one moral community is permitted (or perhaps even required) in another. There is no inconsistency between an agent in the first community saying “ $\phi$ -ing is wrong” and an agent from the other other saying “ $\phi$ -ing is not wrong”, but the prescriptions offered do differ. In those cases it is difficult to see what determines which commands are correct in which contexts. Scanlon notes that if, as is the case for many relativists, the correct set of moral standards is determined by the society, relativism may be seen as a...

“...debunking doctrine, according to which morality is is *merely* a matter of social convention – where the ‘merely’ reflects the assumption that being generally accepted in a society could not, by itself, confer anything like the authority that moral judgements are commonly supposed to have.”

(Scanlon, 1998, p.333)

Here the relativist is left with another dilemma. They must either [1] accept the debunking, and concede that morality does not have the moral authority typically associated with it, or [2] give an explanation why the moral judgements can still confer that authority. Again, I opt for the second of these two routes.

Scanlon describes types of relativism which aim to preserve the moral authority of morality – those rejecting the debunking claims, even if they accept that the moral standards in play may depend on social convention – as varieties of ‘benign relativism’. The position I am advocating is of this sort.

### **Responding to Relativistic Worries**

Turning to the complaint that a relativistic approach delegitimises moral condemnation. There are three types of cases that might occur where it

could seem appropriate to morally condemn someone.<sup>49</sup> It could be someone in our moral community, doing something deemed impermissible in our moral community. Secondly, someone outside of our moral community could be doing something deemed impermissible in *both* our and their moral communities. Thirdly, we could have someone outside of our moral community doing something judged impermissible in our moral community, but accepted as permissible in theirs. In the first two cases, there seems to be no problem. Both agents accept that the conduct is unacceptable, so both could accept that condemnation might be appropriate.

The third case does provide a case where reasonable agents will disagree about the legitimacy of condemnation. If there is no objectively privileged method of discerning what obligations people have – rather, that there are many rival ways such verdicts might be reached – this type of case, of genuine moral disagreement about obligations, is to be expected. Agents may appeal to considerations of various values, as presumably there will some common ground in the acceptance of certain good-making features, in order to engage in a second-order moral discourse. For example, consider a conflict between two communities which see the same things as valuable, the first of which accepts a minimal picture of obligations and the second which has a significantly more burdensome picture. Each could point to values which are better promoted under their framework – perhaps freedom from infringement of morality in the first, and typical aggregate welfare in the second. If the communities do recognise the values pointed to by each other, they will see some legitimacy to the claims of the other, as the Reasons Weighting of the other *would* better promote some value.

A more difficult situation would obtain if the goods that one community sees as having a Reasons Weighting are not acknowledged as valuable at all by another community. In these situations, communities will not even really be able to communicate their moral reasons; they may understand the ‘rules’ but may see them as bizarre or fetishistic.

When there is sufficient common ground, moral discourse can take place. Each party can attempt to convince the other to change, adopting a more (or less) demanding Reasons Weighting system, but so long as both parties

---

<sup>49</sup>Of course, there will be cases where people condemn others inappropriately, i.e. when it would *not* be reasonable to require them to behave otherwise, but this does not seem to offer any special problem.

have a consistent picture of obligation – one that meets the standards of reasonableness – each can only appreciate reasons given by their perceptions of value.

With regards to Scanlon's third problem for relativism, that it undermines the authority of morality, I do not think this is as grave an issue as Scanlon suggests. At the heart of this objection seems to be the claim that morality is not taken *seriously* enough. Morality has, so the objection to relativism goes, an authority that we are clearly not giving proper due if we think it so restricted (to various moral communities) or flimsy (given that what one judges as required in a situation could have been so different). Contrary to this claim, I think the *opponents* are actually the ones failing to give morality its due. If the extremists (or even Scanlon, apparently, given how demanding his contractualism seems to be) are correct, the vast majority of people are rampantly flouting their moral obligations, mostly without concern. It seems puzzling to assert that morality has ultimate authority, that we accept that it has this authority, and yet agents flouting morality's requirements are ubiquitous. However, an account of the sort I have provided can give morality serious authority over an agent's actions, because the agent *identifies* with the standard of reasonableness they accept. In a vaguely Kantian move, we might suggest that the authority of morality comes from the notion that we impose it upon ourselves.

In addition to the reasons for resisting relativism Scanlon provides, we might also suggest that the correctness of relativism presumes a widespread misunderstanding of morality. It is often supposed that claims about what an agent is required to do are universal in some important sense. If the proposal I have offered is accurate, this belief is fundamentally mistaken; we are systematically wrong about what our claims and judgements about moral obligation actually mean. Is this a problem?

Prima facie, this does appear to be a difficulty. The purported generality of morality is something that seems to have widespread acceptance. Most proponents of moral theories take themselves to be saying something universal. This is clear in the work of Kant, with his taking generality to be a requirement of moral laws, but this presumption is clear in any attempts to provide a theory of obligation.

I would reply that some similar error is required in most, if not all, competing views. As was discussed in chapter 1, most plausible theories of

obligation seem susceptible to demandingness objections.<sup>50</sup> If any of these theories were true, they would require that agents are significantly (and probably systematically) mistaken about the extent of their moral obligations. When faced with a choice between positions which claim that seemingly well-informed agents are wildly mistaken about the extent of their moral obligations, and a position claiming those same agents are mistaken about *who* falls under the purview of those obligations, it is unclear if one side wins out. At least, for the position I am arguing for, people typically have access to their own obligations, which, given the picture I have presented, is what they would be considerably more familiar with anyway. We might not see it as surprising that – given widespread moral agreement – it is often taken that the normative force of certain reasons is universal.

#### 6.4.4 Success Evaluation

During the previous chapter I introduced some success criteria for theories of moral obligation. Now that this theory has been cashed out somewhat, we can consider how well it is able to respond to these criteria. The criteria arrived at were that a suitable theory should:

1. Conform to most/all strong intuitions
  - a) Provide ‘moderate’ verdicts.
  - b) Support intuitions about moral progress.
  - c) Must be able to satisfactorily respond to Cullity’s iterative demands problem.
2. Allow for moral discourse.
3. Allow for the possibility of moral mistakes.

Moral mistakes are possible in all the ways previously mentioned.<sup>51</sup>

As has been mentioned already, that the account provided will conform with most (or all) of our intuitions is guaranteed by its incorporation of a community’s own notions of reasonableness. Those determine the extent of

---

<sup>50</sup>An obvious exception being Norcross’s scalar utilitarianism, but this seems to require that we make what seems like an even more serious error – mistakenly believing in the existence of moral obligations.

<sup>51</sup>In section 6.4.2.

the moral obligations, and as they are accepted by members of the moral community (definitionally), the intuitions of the members will conform to the obligations.

More specifically, the picture described seems to support a moderate view: our obligations might demand more of agents – perhaps significantly more – than most of us actually do. Many of us do think that, in terms of time spent on beneficent activities and our monetary contributions to the world’s poor, we ought to do a lot better. At a given time, people might rationalise their behaviour – unfairly treat their reasons as granting justifications or excuses, when they would not for others – to mitigate or alleviate feelings of guilt. They may, for example, tell themselves that they *really need* the money right now, exaggerate how much worse their lives would be without some luxury good, or convince themselves they have discharged their obligation in some other way (“I already gave quite a lot last week, so it’s okay not to this time”). The general thought that we are not doing enough does seem commonplace.

One commonly accepted belief is that there has, over human history, been considerable moral progress. We can look back to fairly recent history and see that slavery was rife, even in the supposedly developed countries. While things are by no means perfect now, we accept that there has been moral progress – the moral landscape is better now than it was when the UK and USA were exploiting slaves en masse. Unger contends that the “idea of moral progress” has something we’re greatly influenced by in our typical moral thinking (1996, p.18). A relativistic picture might be seen to struggle with justifying that moral progress can occur. If there are no objectively right or wrong verdicts about our obligations, it could be asked, how can we be said to improve? A typical, non-relativistic picture could suggest that there are objective moral truths about what our obligations are and that getting closer to understanding and fulfilling them is progress.

In response to this, I claim that distinguishing between axiology-talk and obligation-talk can help us. While I have not endorsed a particular axiology here, so long as the picture is supplemented by an axiology where we can accurately compare moral states of affairs<sup>52</sup> it will be possible to talk of progress. If, for example, we accepted a purely hedonic evaluation, we

---

<sup>52</sup>By this, I do not intend anything philosophically deep, but rather some way of evaluating the world morally.

could say that we have exhibited moral progress if the obligations we now accept result in more pleasure. Things would be more complicated with a pluralistic picture, but I see no reason why, in principle, we could not reach a view of moral progress with this account just as easily as in any other account.

Additionally, to explain moral progress, we may point to various mistakes that have been made by previous moral communities. They may have made a variety of mistakes we have since learnt from.<sup>53</sup> Correcting moral mistakes clearly seems like progress. Alternatively, they may have accepted rules that failed to meet the reasonableness constraint, i.e. regarding bodily integrity as very important, but only for certain people. Systems that permit slavery are likely to be unreasonable.<sup>54</sup>

As was discussed previously, Scanlon has a difficulty when people from (what I would term) different moral communities discuss a practice, where both are considering the same reasons, but they arrive at a different conclusion. This problem does not arise the same way with the type of account I have described, but similar worries may result. When one agent from Moral Community A says that practice  $\phi$  is forbidden, and an individual from Moral Community B says  $\phi$  is not forbidden (both earnestly), there are several options for what might be happening. First, it is possible that one of the agents has made a moral mistake. It could be that an individual from either community is ignorant of some fact, rationalising, or not paying sufficient attention, and has arrived at a mistaken verdict. This would be easily characterised as a moral disagreement, where one agent is correct and the other mistaken.

Alternatively, their Reasons Weightings might actually lead them to different verdicts. Perhaps  $\phi$ -ing really is impermissible by the lights of those in Moral Community A, but permissible by those in Moral Community B. From what I have said so far, it is not obvious that we can actually charac-

---

<sup>53</sup>Of the sort discussed in section 6.4.2.

<sup>54</sup>I say *likely to* because depending upon the goods valued by the society, it is possible to allow for slavery to be permitted. Systems like this would be quite alien to us, however, as the grounding for rights of bodily integrity would have to be devised in some strange way, so that the reasons we judge it to be important would need to be accompanied by some factor we would deem arbitrary (like membership of some race or social class) in order to be generated. With regards to the slavery-permitting systems in our actual history, however, it *does* seem like either the system must have been unreasonable or some mistakes in application were made.

terise this as a disagreement. To elucidate, we could describe  $\phi$  as  $\text{wrong}_A$  but  $\text{permissible}_B$ , which, in itself, involves no disagreement. In some cases this might be what is going on. The agents might simply be talking past each other.

In most of these cases, however, a disagreement (or something like a disagreement) does seem to be happening. There are a wide range of options for explaining this available. To explore them properly would require far more space than I have here, but here are some suggestions.

We might think that when an agent is engaged in moral talk with another they endorse some imperative in such a way that implies that the audience *has a reason* to also accept that imperative.<sup>55</sup> When a person is trying to convince another to change their perspective on a moral issue – trying to convince a meat-eater that eating meat is wrong, for example – there are a number of ways this could be explained.

First, they may think that the other person has a sufficiently similar Reasons Weighting but is making a moral mistake. They might think that presenting the issue in a slightly different way could reveal a tension or inconsistency in the interlocutor’s position and convince them that way. A moral mistake could also easily result from a factual one. Consider two people disagreeing about the permissibility of eating fish because they disagree about whether fish can feel pain. Second, they might accept that the other person does not currently adopt a Reasons Weighting that forbids meat-eating, but want to convince them that they should adopt a new Reasons Weighting. Some of this convincing might require alerting to axiological considerations. Praising and blaming practices may also cause people to make some sort of moral shift. This might be particularly successful if they think the other person is rationalising, as it would inform them that their deception (or self-deception) is noted and subject to some form of sanction. Third, they might be mistaken about what is entailed by their interlocutor’s notion of reasonableness, thinking they are pointing out a mistake, when actually they are not.

While I have not been able to give a fully detailed account here – doing so would require an entire new thesis(!) – I have gestured towards some

---

<sup>55</sup>Gilbert Harman makes a similar move with regards to his “inner judgments” which require that a speaker endorse the position held and that they believe that the audience also does (1975, p.8).

ways the schematic could be developed.

#### **6.4.5 A Concession to the Extremists**

Despite these claims that an account of this sort might not be quite as extreme as could be suggested, it certainly contains features that some theorists would be reluctant to endorse – like its relativism. However, even if one is dissuaded by the relativism about moral obligation, it might still be the case that this account offers something appealing. Consider the extremists discussed in 2. These theorists must face the challenge of explaining away many of our strong intuitions. One potential explanation for why intuitions provide some judgements would be given by the account discussed.

Recent psychological literature has suggested that our judgements of moral obligation operate in a way strongly connected to common practice. The “Common is Moral” Heuristic, observed by Lindström et al. holds that when making judgements about the social behaviour of others, people typically “substitute the target attribute “moral value” with the more accessible attribute “behavioral frequency” ” (2017). This means that when considering whether something is morally permissible, an agent will, due to a cognitive association between commonness and moral values, often judge the behaviour acceptable if they judge it to be common.

An extremist may accept that the type of picture I have provided, of obligation-statuses of actions being judged based on whether it seems reasonable (which is in turn informed by what is common), as providing evidence for how our intuitions mislead us. They could take (at least elements of) this account to provide a descriptive model for how people often judge moral obligations. This would leave the extremist with the task of explaining why the moral data they make use of is immune to this type of error.

Of course, I have argued here that the demandingness objection is a good one, and therefore that we can rely on the underlying intuitions. For my dialectic opponents, however, understanding our moral judgements as being formed in a way similar to the present proposal may prove fruitful.

# Bibliography

- van Ackeren, Marcel and Sticker, Martin (2015): 'Kant and Moral Demandingness,' *Ethical Theory and Moral Practice*, 18(1), pp. 75–89.
- Anscombe, G. E. M. (1958): 'Modern Moral Philosophy,' *Philosophy*, 33(124), pp. 1–19.
- Appiah, Kwame Anthony (2007): *Cosmopolitanism: Ethics in a World of Strangers*, London: Penguin.
- Ashford, Elizabeth (2000): 'Utilitarianism, Integrity, and Partiality,' *Journal of Philosophy*, 97(8), pp. 421–439.
- Ashford, Elizabeth (2003): 'The Demandingness of Scanlon's Contractualism,' *Ethics*, 113, pp. 273–302.
- Ashford, Elizabeth and Mulgan, Tim (2012): 'Contractualism,' in Edward N. Zalta, (ed.) *The Stanford Encyclopedia of Philosophy*, Metaphysics Research Lab, Stanford University, fall 2012 edition.  
**URL:** <http://plato.stanford.edu/archives/fall2012/entries/contractualism/>
- Austin, J. L. (1961): 'A Plea for Excuses,' in J. O. Urmson and G. J. Warnock, (eds.) *Austin, Philosophical Papers*, pp. 123–154, Oxford: Oxford University Press.
- Baron, Marcia (1987): 'Kantian Ethics and Supererogation,' *The Journal of Philosophy*, 84(5), pp. 237–262.
- Baron, Marcia (1995): *Kantian Ethics Almost Without Apology*, Ithaca: Cornell University Press.
- Baron, Marcia (2005): 'Justifications and Excuses,' *Ohio State Journal of Criminal Law*, 2, pp. 387–406.

- Baron, Marcia (2007): 'Excuses, Excuses,' *Criminal Law and Philosophy*, 1(1), pp. 21–39.
- Baron, Marcia (2011): 'The Standard of the Reasonable Person in the Criminal Law,' in R.A. Duff, Lindsay Farmer, S.E. Marshall, Massimo Renzo and Victor Tadros, (eds.) *The Structures of the Criminal Law*, pp. 11–35, Oxford: Oxford University Press.
- Baron, Marcia (2017): 'Reasonableness,' Unpublished manuscript.
- Barry, Brian (1989): *Theories of Justice*, California: University of California Press.
- Bell, Macalester (2013): 'The Standing to Blame: A Critique,' in D. Coates and N. Tognazzini, (eds.) *Blame: Its Nature and Norms*, pp. 263–281, Oxford University Press.
- Benn, Claire (2016): 'Over-demandingness Objections and Supererogation,' in M. van Ackeren and M. Kuhler, (eds.) *The Limits of Moral Obligation: Moral Demandingness and Ought Implies Can*, New York: Routledge.
- Berlin, Isaiah (2002): *Liberty*, Oxford: Oxford University Press.
- Blum, Lawrence (1998): 'Community and Virtue,' in Roger Crisp, (ed.) *How Should One Live: Essays on the Virtues*, pp. 231–250, 1998.
- Bourbaki, Nicola (2001): 'Living High and Letting Die,' *Philosophy*, 76(297), pp. 435–442.
- Braddock, Matthew (2013): 'Defusing the Demandingness Objection: Unreliable Intuitions,' *Journal of Social Philosophy*, 44(2), pp. 169–191.
- Brink, David (1997): 'Self-Love and Altruism,' *Social Policy and Philosophy*, 14(1), pp. 122–157.
- Cappelen, Herman (2012): *Philosophy Without Intuitions*, Oxford: Oxford University Press.
- Carbonell, Vanessa (2012): 'The Ratcheting-Up Effect,' *Pacific Philosophical Quarterly*, 93, pp. 228–254.

- Carbonell, Vanessa (2015): 'Review: Moral Failure: On the Impossible Demands of Morality,' *Notre Dame Philosophical Reviews*.  
**URL:** <http://ndpr.nd.edu/news/moral-failure-on-the-impossible-demands-of-morality/>
- Carbonell, Vanessa (2016): 'Differential Demands,' in M. van Ackeren and M. Kuhler, (eds.) *The Limits of Moral Obligation: Moral Demandingness and Ought Implies Can*, pp. 36–50, New York: Routledge.
- Chappell, Richard Yetter (2017): 'Willpower Satisficing,' *Noûs*.  
**URL:** <http://dx.doi.org/10.1111/nous.12213>
- Chappell, Timothy (2007): 'Integrity and Demandingness,' *Ethical Theory and Moral Practice*, 10(3), pp. 255–265.
- Cohen, G. A. (2000a): 'If You're an Egalitarian, How Come You're so Rich?' *The Journal of Ethics*, 4(1/2), pp. 1–26.
- Cohen, G. A. (2000b): *Karl Marx's Theory of History: A Defence (Expanded Edition)*, Princeton: Princeton University Press.
- Conly, Sarah (1985): 'The Objectivity of Morals and the Subjectivity of Agents,' *American Philosophical Quarterly*, 22(4), pp. 275–286.
- Cullity, Garrett (2004): *The Moral Demands of Affluence*, Oxford: Clarendon Press.
- Cullity, Garrett (2009): 'Demandingness Arguments from Presupposition,' in Timothy Chappell, (ed.) *The Problem of Moral Demandingness*, pp. 8–34, Basingstoke: Palgrave MacMillan.
- Dorsey, Dale (2013): 'The Supererogatory and How to Accommodate It,' *Utilitas*, 25(3), pp. 355–382.
- Dressler, Joshua (1995): *Understanding Criminal Law*, United States of America: Times Mirror, second edition.
- Driver, Julia (1992): 'The Suberogatory,' *Australasian Journal of Philosophy*, 70(3), pp. 286–295.
- Fishkin, James (1982): *The Limits of Obligation*, London: Yale University Press.

- Foot, Philippa (1967): ‘The Problem of Abortion and the Doctrine of Double Effect,’ *Oxford Review*, 5, pp. 1–7.
- Frankfurt, Harry (1998a): ‘Alternative Possibilities and Moral Responsibility,’ in *The Importance of What We Care About*, chapter 1, pp. 1–10, New York: Cambridge University Press.
- Frankfurt, Harry (1998b): ‘The Problem of Action,’ in *The Importance of What We Care About*, chapter 6, pp. 69–79, New York: Cambridge University Press.
- Garvey, Stephen (2009): ‘When Should a Mistake of Fact Excuse?’ *Texas Tech Law Review*, 42, pp. 359–382.
- Gert, Bernard (2004): *Common Morality: Deciding What to Do*, Oxford: Oxford University Press.
- Gibbard, Alan (1990): *Wise Choices, Apt Feelings: A Theory of Normative Judgement*, Oxford: Clarendon Press.
- Goodin, Robert (2009): ‘Demandingness as a Virtue,’ *The Journal of Ethics*, 13(1), pp. 1–13.
- Gustafsson, Johan (2016): ‘Consequentialism with Wrongness Depending on the Difficulty of Doing Better,’ *Thought*, 5, pp. 108–118.
- Haji, Ishtiyaque (1997): ‘An Epistemic Dimension of Blameworthiness,’ *Philosophy and Phenomenological Research*, 57(3), pp. 523–544.
- Hallie, Philip (1979): *Lest Innocent Blood Be Shed*, London: Michael Joseph.
- Hare, Caspar (2013): *The Limits of Kindness*, Oxford: Oxford University Press.
- Harman, Gilbert (1975): ‘Moral Relativism Defended,’ *The Philosophical Review*, 84(1), pp. 3–22.
- Harman, Gilbert and Thomson, Judith Jarvis (1996): *Moral Relativism and Moral Objectivity*, Blackwell.
- Harris, John (1975): ‘The Survival Lottery,’ *Philosophy*, 50(191), pp. 81–87.

- Hedberg, Trevor (2014): 'Epistemic supererogation and its implications,' *Synthese*, 191, pp. 3621–3637.
- Hill, Thomas (2002): *Human Welfare and Moral Worth: Kantian Perspectives*, Oxford: Oxford University Press.
- Hills, Alison (2010): 'Utilitarianism, Contractualism and Demandingness,' *The Philosophical Quarterly*, 60(239), pp. 225–242.
- Hinsliff, Gaby (2017): 'The 'Arm Vagina' – Hollywood's Latest form of female self-flagellation,' *The Guardian*.  
**URL:** <https://www.theguardian.com/commentisfree/2017/dec/01/arm-vagina-hollywood-female-self-flagellation-jennifer-lawrence>
- Hobbes, Thomas (1995): *Leviathan*, Charlottesville: InteLex.
- Hooker, Brad (1990): 'Rule-Consequentialism,' *Mind*, 99(393), pp. 67–77.
- Hooker, Brad (1991): 'Brink, Kagan and Self-Sacrifice,' *Utilitas*, 3(2), pp. 263–273.
- Hooker, Brad (1999): 'Review: Sacrificing for the Good of Strangers – Repeatedly,' *Philosophy and Phenomenological Research*, 59(1), pp. 177–181.
- Hooker, Brad (2000): *Ideal Code, Real World: A Rule-consequentialist Theory of Morality*, Oxford: Clarendon Press.
- Hooker, Brad (2009a): 'The Demandingness Objection,' in Timothy Chappell, (ed.) *The Problem of Moral Demandingness*, pp. 148–162, Basingstoke: Palgrave MacMillan.
- Hooker, Brad (2009b): 'When is Impartiality Morally Appropriate?' in Brian Feltham and John Cottingham, (eds.) *Partiality and Impartiality: Morality, Special Friendships and the Wider World*, pp. 26–41, Oxford: Oxford University Press.
- Hooker, Brad (2016): 'Wrongness, Evolutionary Debunking, Public Rules,' *Ethica & Politca / Ethics & Politics*, XVIII(1), pp. 135–149.
- Hume, David (1969): *Treatise of Human Nature*, Bucks.: Penguin (1790).
- Jacobson, Daniel (2008): 'Utilitarianism without Consequentialism: The Case of John Stuart Mill,' *The Philosophical Review*, 117, pp. 159–191.

- Jonathan, Dancy (2004): *Ethics Without Principles*, Oxford: Oxford University Press.
- Kagan, Shelly (1989): *The Limits of Morality*, Oxford: Oxford University Press.
- Kagan, Shelly (2001): 'Thinking About Cases,' *Social Philosophy & Policy*, 18(2), pp. 44–63.
- Kamm, Frances (1999): 'Review of *Living High and Letting Die* by Peter Unger,' *The Philosophical Review*, 108(2), pp. 300–305.
- Kant, Immanuel (1992): 'Of an Alleged Right to Lie from Philanthropy,' in Mary Gregor, (ed.) *Practical Philosophy*, pp. 605–616, Cambridge University Press.
- Kant, Immanuel (1996): *The Metaphysics of Morals*, Cambridge: Cambridge University Press. Mary Gregor (trans.).
- Kelly, Erin I. (2013): 'What is an Excuse?' in D. Coates and N. Tognazzini, (eds.) *Blame: Its Nature and Norms*, pp. 242–262, Oxford University Press.
- de Lazari-Radek, Katarzyna and Singer, Peter (2010): 'Secrecy in Consequentialism: A Defence in Esoteric Morality,' *Ratio*, XXIII, pp. 34–58.
- Lenman, James (2000): 'Consequentialism and Cluelessness,' *Philosophy and Public Affairs*, 29(4), pp. 342–370.
- Lichtenberg, Judith (2010): 'Positive Duties, Negative Duties and the "New Harms",' *Ethics*, 120(3), pp. 557–578.
- Lindström, B., Jangard, S., Selbing, I. and Olsson, A. (2017): 'The Role of a "Common Is Moral" Heuristic in the Stability and Change of Moral Norms,' *Journal of Experimental Psychology: General*, Online First Publication, pp. 1–15.
- MacFarquhar, Larissa (2015): *Strangers Drowning: Grappling With Impossible Idealism, Drastic Choices, and the Overpowering Urge to Help*, New York: Penguin.

- MacIntyre, Alastair (1984): *Is Patriotism a Virtue? (The Lindley Lecture)*, Lawrence: University of Kansas.
- McElwee, Brian (2007): ‘Consequentialism, Demandingness and the Monism of Practical Reason,’ *Proceedings of the Aristotelian Society, New Series*, 107, pp. 359–374.
- McElwee, Brian (2010): ‘The rights and wrongs of consequentialism,’ *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 151(3), pp. 393–412.
- McElwee, Brian (2015): ‘The Value of Virtues,’ *Utilitas*, 27, pp. 61–81.
- McElwee, Brian (2016a): ‘Supererogation Across Normative Domains,’ *Australasian Journal of Philosophy*.
- McElwee, Brian (2016b): ‘What is Demandingness?’ in M. van Ackeren and M. Kuhler, (eds.) *The Limits of Moral Obligation: Moral Demandingness and Ought Implies Can*, pp. 19–35, New York: Routledge.
- McElwee, Brian (2017): ‘Demandingness Objections in Ethics,’ *Philosophical Quarterly*, 67, pp. 84–105.
- McGeer, Victoria (2013): ‘Civilizing Blame,’ in D. Coates and N. Tognazzini, (eds.) *Blame: Its Nature and Norms*, pp. 162–188, Oxford University Press.
- McNamara, Paul (2014): ‘Deontic Logic,’ in Edward N. Zalta, (ed.) *The Stanford Encyclopedia of Philosophy*, Metaphysics Research Lab, Stanford University, winter 2014 edition.
- Mele, Alfred (2003): *Motivation and Agency*, Oxford: Oxford University Press.
- Metz, Thaddeus (2001): ‘Review of *Moral Demands in Nonideal Theory*,’ *The Philosophical Review*, pp. 614–617.
- Mill, John Stuart (1861): ‘*Utilitarianism*,’ in pp.61-114 of H. West, (ed.) *The Blackwell Guide to Mill’s Utilitarianism*, (2006), Oxford: Blackwell.
- Mill, John Stuart (1865): *Auguste Comte and Positivism*, London: Trubner & Co.

- Miller, Richard (2004): 'Beneficence, Duty and Distance,' *Philosophy & Public Affairs*, 32(4), pp. 357–383.
- Miller, Richard (2010): *Globalizing Justice: The Ethics of Poverty and Power*, Oxford: Oxford University Press.
- Moore, Margaret (1996): 'On Reasonableness,' *Journal of Applied Philosophy*, 13(2), pp. 167–178.
- Mulgan, Tim (2000): 'Review: *Living High and Letting Die*,' *Mind*, 109(434), pp. 397–400.
- Mulgan, Tim (2001): *The Demands of Consequentialism*, Oxford: Clarendon Press.
- Mulgan, Tim (2003): 'Review: Liam Murphy's *Moral Demands in Nonideal Theory*,' *Utilitas*, pp. 113–116.
- Mulgan, Tim (2006): *Future People: A Moderate Consequentialist Account of our Obligations to Future Generations*, Oxford: Oxford University Press.
- Murphy, Liam (1993): 'The Demands of Beneficence,' *Philosophy & Public Affairs*, 22(4), pp. 267–292.
- Murphy, Liam (2000): *Moral Demands in Nonideal Theory*, New York: Oxford University Press.
- Nagel, Thomas (1986): *The View From Nowhere*, Oxford: Oxford University Press.
- Nagel, Thomas (1991): *Equality and Partiality*, Oxford: Oxford University Press.
- Narveson, Jan (2003): 'We Don't Owe Them a Thing! A Tough-minded but Soft-hearted View of Aid to the Faraway Needy,' *The Monist*, 86(3).
- Norcross, Alastair (1997): 'Good and Bad Actions,' *The Philosophical Review*, 106(1), pp. 1–34.
- Norcross, Alastair (2006): 'The Scalar Approach to Consequentialism,' in H. West, (ed.) *The Blackwell Guide to Mill's Utilitarianism*, Oxford: Blackwell.

- Nozick, Robert (1974): *Anarchy, State and Utopia*, New York: Basic Books.
- Parfit, Derek (1984): *Reasons and Persons*, Oxford: Clarendon Press.
- Parfit, Derek (2011): *On What Matters: Volume 1*, Oxford: Oxford University Press.
- Pinheiro Walla, Alice (2015): ‘Kant’s Moral Theory and Demandingness,’ *Ethical Theory and Moral Practice*, 18(4).
- Price, Richard (1789): *A Discourse on the Love of our Country*, London: The Strand.
- Pybus, Elizabeth (1982): ‘Saints and Heroes,’ *Philosophy*, 57(220), pp. 193–199.
- Railton, Peter (1984): ‘Alienation, Consequentialism and the Demands of Morality,’ *Philosophy & Public Affairs*, 13(2), pp. 134–171.
- Rawls, John (1991): *A Theory of Justice*, Massachusetts: Belknap Press (1971).
- Rawls, John (1993): *Political Liberalism*, New York: Columbia University Press.
- Rescher, Nicholas (1954): ‘Reasonableness in Ethics,’ *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 5(4), pp. 58–62.
- Ridge, Michael (2009): ‘Fairness and Non-Compliance,’ in Brian Feltham and John Cottingham, (eds.) *Partiality and Impartiality: Morality, Special Friendships and the Wider World*, pp. 194–222, Oxford: Oxford University Press.
- Rohwetter, Marcus Von (2017): ‘Wie ich als Verbraucher beinahe den Verstand verlor,’ *Zeit*.  
**URL:** <http://www.zeit.de/2017/42/konsum-verbraucher-verantwortungsbewusstsein-waren>
- Ross, W. D. (2002): *The Right and The Good*, Oxford: Oxford University Press.

- Scanlon, T. M. (1998): *What We Owe to Each Other*, London: Belknap Press.
- Scanlon, T. M. (2008): *Moral Dimensions: Permissibility, Meaning, Blame*, Cambridge: Mass.: Harvard University Press.
- Scanlon, T. M. (2013): 'Interpreting Blame,' in D. Coates and N. Tognazzini, (eds.) *Blame: Its Nature and Norms*, pp. 84–99, Oxford University Press.
- Scanlon, T. M. (2014): *Being Realistic About Reasons*, Oxford: Oxford University Press.
- Scheffler, Samuel (1982): *The Rejection of Consequentialism*, Oxford: Oxford University Press.
- Scheffler, Samuel (1992): *Human Morality*, Oxford: Oxford University Press.
- Scheffler, Samuel (1994): 'Prerogatives without Restrictions,' in *The Rejection of Consequentialism: Revised Edition*, pp. 167–189, Oxford: Clarendon Press.
- Sher, George (2006): *In Praise of Blame*, Oxford: Oxford University Press.
- Sher, George (2013): 'Wrongdoing and Relationships: The Problem of the Stranger,' in D. Coates and N. Tognazzini, (eds.) *Blame: Its Nature and Norms*, pp. 49–65, Oxford University Press.
- Sibley, W. B. (1953): 'The Rational Versus the Reasonable,' *Philosophical Review*, 62, pp. 554–60.
- Sidgwick, Henry (1907): *Methods of Ethics*, London: Macmillan and co, seventh edition edition.
- Singer, Peter (1972): 'Famine, Affluence & Morality,' *Philosophy & Public Affairs*, 1(3), pp. 229–243.
- Singer, Peter (1993): *Practical Ethics*, Cambridge: Cambridge University Press, second edition edition.
- Singer, Peter (2007): 'Review: Review Essay on "The Moral Demands of Affluence",' *Philosophy and Phenomenological Research*, 75(2), pp. 475–483.

- Singer, Peter (2010): *The Life You Can Save*, New York: Random House.
- Singer, Peter (2015): *Famine, Affluence and Morality*, Oxford: Oxford University Press.
- Singer, Peter and de Lazari-Radek, Katarzyna (2014): *The Point of View of the Universe: Sidgwick and Contemporary Ethics*, Oxford: Oxford University Press.
- Singer, Peter and de Lazari-Radek, Katarzyna (2016): ‘Doing our Best for Hedonistic Utilitarianism,’ *Ethics & Politics*, 18(1), pp. 187–207.
- Skorupski, John (2011): *The Domain of Reasons*, Oxford: Oxford University Press.
- Slote, Michael (1984): ‘Satisficing Consequentialism,’ *Proceedings of the Aristotelian Society, Supplementary Volumes*, 58, pp. 139–163.
- Smart, J. J. C. (1961): ‘Free-will, praise and blame,’ *Mind*, 70(279), pp. 291–306.
- Smith, Holly (1983): ‘Culpable Ignorance,’ *The Philosophical Review*, 92(4), pp. 543–571.
- Sobel, David (2007): ‘The Impotence of the Demandingness Objection,’ *Philosophers’ Imprint*, 7(8), pp. 1–17.
- Strawson, P. F. (2008): *Freedom and Resentment and other essays*, London: Routledge.
- Swanton, Christine (2009): ‘Virtue Ethics and the Problem of Demandingness,’ in Timothy Chappell, (ed.) *The Problem of Moral Demandingness*, pp. 104–122, Basingstoke: Palgrave MacMillan.
- Tanswell, Fenner Stanley (2017): ‘Conceptual engineering for mathematical concepts,’ *Inquiry*, 0(0), pp. 1–33.  
**URL:** <https://doi.org/10.1080/0020174X.2017.1385526>
- Taurek, John (1977): ‘Should the Numbers Count?’ *Philosophy & Public Affairs*, 6(4), pp. 293–316.
- Tessman, Lisa (2015): *Moral Failure: On the Impossible Demands of Morality*, Oxford: Oxford University Press.

- Thomson, Judith Jarvis (1971): ‘A Defense of Abortion,’ *Philosophy & Public Affairs*, 1(1), pp. 47–66.
- Timmerman, Travis (2015): ‘Sometimes there is nothing wrong with letting a child drown,’ *Analysis*, 75(4), pp. 204–212.
- Timmermann, Jens (2005): ‘Good but Not Required? - Assessing the Demands of Kantian Ethics,’ *Journal of Moral Philosophy*, 2(1), pp. 9–27.
- Timmermann, Jens (forthcoming): ‘Autonomy, Progress and Virtue: Why Kant Has Nothing to Fear from the Overdemandingness Objection,’ *Kantian Review*.
- Unger, Peter (1996): *Living High and Letting Die: Our Illusion of Innocence*, Oxford: Oxford University Press.
- Urmson, J. O. (1958): ‘Saints and Heroes,’ in A. I. Melden, (ed.) *Essays in Moral Philosophy*, pp. 119–216, University of Washington Press, London.
- Waismann, F (1968): ‘Verifiability,’ in A. Flew, (ed.) *Logic and Language*, p. 118–144, Oxford: Basil Blackwell.
- Wenar, Leif (2010): ‘Poverty is No Pond: Challenges for the Affluent,’ in T. Illingworth, P. Pogge and L. Wenar, (eds.) *Giving Well: The Ethics of Philanthropy*, pp. 104–132, Oxford: Oxford University Press.
- Williams, Bernard (1981): *Moral Luck*, New York: Cambridge University Press.
- Williams, Bernard and Smart, J.C.C. (1973): *Utilitarianism: For and Against*, Cambridge: Cambridge University Press.
- Woollard, Fiona (2016a): ‘Dimensions of Demandingness,’ *Proceedings of the Aristotelian Society*, 116(1), pp. 89–106.
- Woollard, Fiona (2016b): ‘Motherhood and Mistakes about Defeasible Duties to Benefit,’ *Philosophy and Phenomenological Research*.  
**URL:** <http://dx.doi.org/10.1111/phpr.12355>
- Zagzebski, Linda (2017): *Exemplarist Moral Theory*, Oxford: Oxford University Press.