

The London School of Economics and Political Science

Imperfect Duties in Current Debates
Supererogation, Demandingness, and
Collective Impact Cases

Dmitry Ananiev

A thesis submitted to the Department of Philosophy, Logic and Scientific Method
of the London School of Economics and Political Science
for the degree of Doctor of Philosophy, 29th March 2023

Declaration

I certify that the thesis I have presented for examination for the MPhil/PhD degree of the London School of Economics and Political Science is solely my own work.

The copyright of this thesis rests with the author. Quotation from it is permitted, provided that full acknowledgement is made. This thesis may not be reproduced without my prior written consent. I warrant that this authorisation does not, to the best of my belief, infringe the rights of any third party.

I declare that my thesis consists of 50077 words.

Dmitry Ananiev

Abstract

The concept of imperfect duty has become increasingly prominent in a wide range of debates in moral philosophy. However, many problems related to imperfect duties remain underdiscussed. This thesis aims to fill some of these gaps. Specifically, I focus on three issues.

One is the compatibility of imperfect duties with supererogation. It has been argued that due to their limitless nature, if we accept imperfect duties in our moral theory, then we should abandon the category of supererogation. Contrary to this, I argue that imperfect duties are compatible with supererogation.

The second issue is the connection between latitude permitted by imperfect duties and demandingness of morality. It is often claimed that this latitude implies that morality is moderately demanding. I argue, however, that there is no interesting connection between latitude and demandingness: the implication is true only if we already assume, on independent grounds, that morality is not overly demanding.

The third issue is that imperfect duties seem to help us solve the ‘inefficacy problem’. I argue that imperfect duties fail in this respect. They fail primarily because judgments about one’s compliance with the requirements of an imperfect duty can only be made by assessing one’s life overall. However, on such a scope, it is overwhelmingly unlikely that agents who act in collectively harmful ways violate any imperfect duties in virtue of so acting. But the fact that imperfect duties do not solve the inefficacy problem might not be too regrettable.

To demonstrate this, in the final part of my thesis, I discuss the nature of cases that generate the inefficacy problem. It is thought that in some such cases, individual agents make no moral difference whatsoever to the outcome. I argue that when the harms in question are purely experiential, there are good reasons to think that cases of this kind are impossible. Individual agents in such cases have at least a chance of making a morally significant difference.

Acknowledgements

I am very lucky to have had four wonderful people and excellent philosophers as my supervisors during my time at the LSE. I started my PhD with Alex Voorhoeve as my main supervisor, and I would like to thank him first. My ideas for the main topic of this dissertation first emerged from conversations with Alex. Although he no longer officially supervised my project after my first year, I still had his support whenever I needed it.

Mike Otsuka and Johanna Thoma have been my supervisors throughout most of my PhD journey, and I cannot thank them enough. I am grateful for all the attention they have given to my work, for all the immensely helpful conversations, and for carefully reading through many poor drafts, quite a few of which did not end up in the dissertation (to everyone's benefit). Among many other things, I am grateful to Mike for being exceptionally supportive during the challenging time of the pandemic, and to Johanna for seamlessly taking over as the main supervisor and making the final year of my PhD as stress-free as possible.

I am extremely grateful to Jonathan Parry. Although he officially supervised me for less than a year, he had a tremendous impact on my work. His generous, detailed, and quickly provided feedback pushed me to produce better and clearer arguments. Without his support, my dissertation would have turned out much worse.

I am grateful to all other academics and the professional services staff at the LSE Philosophy Department. You make it a special place, and I have been very happy to work on my dissertation here. I thank Susanne Burri, Richard Bradley, Liam Kofi Bright, Anna Mahtani, and Christian List for their positive influence on my work, be it through written comments, through a conversation, or through a remark at a seminar (sometimes even a small remark has a profound and lasting impact). I cannot talk about the Department without mentioning my fellow PhD students. Fabian, Sophie, Cecily, Charles, Nick, Nic, Margherita, Lorenzo, Eva, and the rest — I thank you for all

the philosophical and non-philosophical discussions, for inspiring me to work harder, and for all the fun and revelry.

I presented parts of my dissertation at several events. I am especially grateful for all the feedback and discussions to the audiences at the LSE Choice Group, “Debunking Inefficacy Reasoning” Workshop in Amsterdam, the Formal Ethics Online Conference, the 2nd Munich Graduate Conference in Ethics, Brave New World conferences organised by MANCEPT, 3rd POLEMO Symposium in Vienna, and HKU Legal and Political Philosophy Graduate Conference in Hong Kong.

My friends outside academia, scattered around the world, have made these past several years better in many ways, and for that I am grateful. Special thanks to Anton, Evgeny, Maxim, and Will for keeping me company regardless of where we all are.

I thank my brothers, Nikolay and Ivan, for always having my back and for visiting me from time to time, providing much needed breaks from work. Ivan’s customary question to me for the past several years has been, “So, have you managed to write anything at all?” I can finally answer, “I suppose I have.”

I thank Alina for always being close and for never doubting that I would get through the PhD even when I had my doubts.

I owe the deepest gratitude to my parents, Alexey and Daria. Without their love, support, and encouragement, I would never have become who I am, would never have got where I am, and would never have completed a project such as this.

Contents

1 INTRODUCTION	9
1.1 Does It Make Sense to Ask What Imperfect Duties Are?.....	10
1.2 Kantian End-Based Views on Imperfect Duty	11
1.3 Kantian End-Based Views in the Context of Current Debates	16
1.3.1 Imperfect Duties and Supererogation.....	16
1.3.2 Imperfect Duties and Latitude	17
1.3.3 Imperfect Duties and Collective Impact Cases	18
1.4 Concluding Remarks	20
2 KANTIAN IMPERFECT DUTIES AND SUPEREROGATION: A CASE FOR COMPATIBILITY	22
2.1 Introduction.....	22
2.2 Imperfect Duties and Supererogation: The Relevant Concepts	24
2.2.1 Kantian Imperfect Duties.....	24
2.2.2 Supererogation	25
2.3 Why Think That Imperfect Duties and Supererogation Are Incompatible? .	27
2.4 Limitlessness and Supererogation: Towards Compatibility	31
2.4.1 Persistence.....	31
2.4.2 Open-Endedness	36
2.4.3 Requiredness	38
2.5 Further Steps: Importance of Weak Supererogation.....	43
2.6 Objections	46
2.6.1 LAR and Weak Supererogation.....	46
2.6.2 Mere Compatibility?	49
2.7 Conclusion	50
3 IMPERFECT DUTIES, LATITUDE, AND DEMANDINGNESS: SEARCHING FOR THE RELATION AND CLARIFYING THE MISCONCEPTIONS.....	52

3.1	Introduction.....	52
3.2	Demandingness and Latitude.....	54
3.2.1	Demandingness	54
3.2.2	Latitude.....	55
3.3	The Two Interpretations of the Relation.....	58
3.4	Against CON-NC and EXP-NC	61
3.4.1	When CON-NC is False	61
3.4.2	When CON-NC is Trivially True	64
3.4.3	EXP-NC: False or Question-Begging	67
3.5	Against CON-C and EXP-C.....	69
3.5.1	When CON-C is False	69
3.5.2	When CON-C is Trivially True	71
3.5.3	EXP-C: False or Question-Begging	72
3.6	How Should the Relation Between Latitude and Demandingness be Construed?	73
3.7	‘Confinement’ as Demandingness?	74
3.8	Conclusion	81
4	IMPERFECT DUTIES AND COLLECTIVE IMPACT CASES	82
4.1	Introduction.....	82
4.2	The Desiderata for a Solution.....	85
4.2.1	Identifying the Desiderata.....	85
4.2.2	Clarifying the Desiderata.....	88
4.3	Kantian Imperfect Duties as a Basis for a Solution	90
4.4	Why Kantian Imperfect Duties Are Unfit as a Basis for a Solution	95
4.4.1	Obligatory Ends and Isolated Patterns of Activity: The Problem of Scope of Assessment	95
4.4.2	Why Imperfect Duties Do Not Help to Solve the Inefficacy Problem	97
4.4.3	Two Reasons Why Making a Difference Wouldn’t Save The Imperfect- Duty-Based Solution.....	100

4.4.4	What We Are Left With	103
4.5	Albertzart's Proposal	105
4.6	Searching for the Contradiction	109
4.6.1	When Contradictions of the Relevant Sort are Virtually Impossible	111
4.6.2	When Contradictions of the Relevant Sort are Rare	115
4.7	Conclusion	118
5 THERE ARE NO (PURELY PHENOMENAL) COLLECTIVE IMPACT		
CASES IN WHICH YOU CAN'T MAKE A DIFFERENCE		119
5.1	Introduction.....	119
5.2	Kinds of Collective Impact Cases: Setting the Focus	120
5.3	Getting to the Paradox.....	124
5.4	Defending the NTPCs While Avoiding the Paradox.....	126
5.5	The Case Against NTPCs: Conceptual Work.....	130
5.5.1	Transitivity Defended.....	130
5.5.2	Instability and Transitivity Without Paradox.....	135
5.5.3	Instability and Transitivity Against the Possibility of NTPCs	136
5.5.4	Instability and Transitivity Against the Prevalence of NTPCs	139
5.5.5	Taking Stock.....	140
5.6	The Case Against NTPCs: Empirical Work	141
5.7	Conclusion	144
6 CONCLUDING REMARKS		146
BIBLIOGRAPHY		151

Chapter 1

Introduction

Does morality only require performing or omitting certain kinds of actions? Or does it also issue broader requirements regarding how one should structure one's life and what goals to pursue? And if it does issue such broader requirements, what are their features and how can they be characterised? These are the main questions motivating this dissertation. I assume that the first question should be answered negatively, but the second one — affirmatively. There are such broader requirements, and, in continuity with the Kantian tradition, I shall refer to them as “imperfect duties”, as opposed to the more familiar requirements on action, which I label “perfect duties”. The third question has received ample attention in the literature and has become increasingly popular over the past several decades.¹ Still, it is far from having a comprehensive answer.² The aim of this dissertation is to contribute to answering this question by addressing some of the problems with characterising imperfect duties and their place in our moral theorising.

In this introduction, I shall first clarify the main question. Second, I shall outline the broadly Kantian view on imperfect duties that I adopt throughout my discussion in this dissertation. Finally, I shall situate this view in the context of the current

¹ Here are just a few examples of recent work on imperfect duties: Greenspan (2010), Hanser (2014), Herman (2021), Portmore (*forthcoming*), Rainbolt (2000), Salam (2014), Schroeder (2014), Segal (2022).

² It is telling that many existing discussions of imperfect duties start by remarking how confusing this concept is or how difficult it is to make sense of the various debates surrounding this concept. See, e.g., Herman (2012: 392–393), Hope (2014: 396), Schroeder (2014: 557), Schumaker (1992: 3–7), Wallace (2019: 50).

debates surrounding imperfect duties. In the process, I shall also give an overview of each chapter, as each chapter is dedicated to a specific problem within these debates.

1.1 Does It Make Sense to Ask What Imperfect Duties Are?

The problem that cuts across many discussions involving imperfect duties is that it is still a matter of debate what imperfect duties are.³ The question regarding what they are can be understood in at least two ways. On one way, the question asks us to try to determine how we should apply the technical terms “perfect duty” and “imperfect duty”. I do not find this to be a promising way of interpreting the question. These terms can be used to pick out different distinctions and using them in different ways does not seem to lead to interesting disagreements. For instance, on one outdated use, these terms refer to the distinction between duties, whose fulfilment is necessary for the existence of society and duties, whose fulfilment improves our lives but is not necessary for the bare existence of society.⁴ On one very common current use, they are used to mark the distinction between duties permitting little or no latitude regarding how to fulfil their requirements and duties permitting significant latitude of that kind.⁵ But many of the distinctions where the distinguished concepts have received the labels of “perfect duty” and “imperfect duty” are theoretically interesting for different reasons. There is little point in trying to argue that only one of these uses picks out *the* distinction.

There is a second, more interesting way of understanding the question. It is widely accepted that the kind of duties that are commonly referred to as imperfect, includes duties such as beneficence, gratitude, sympathy, and perhaps some others. So, the

³ Schroeder (2014), Hope (2014), Herman (2021).

⁴ This view is attributed to early modern philosophers such as Pufendorf. See, e.g., Schneewind (1990: 49–50).

⁵ See Rainbolt (2000) and Hope (2014: 397) for overviews.

second question starts with a conjecture that there is a certain class of duties that share one or more of a list of features. The question, then, is about what specific features unite these duties, why *these* duties share *these* features, and what follows from including these duties in moral theorising. Each of these aspects of the question gives rise to interesting disagreements. Furthermore, addressing this question helps differentiate imperfect duties from perfect duties. It is widely thought that the duties against harming or killing others, or the duty to keep one's promises are perfect. Thinking about the key distinctive features of imperfect duties also helps determine what exactly unites perfect duties.

1.2 Kantian End-Based Views on Imperfect Duty

My dissertation aims to contribute to the debates surrounding some of the disagreements regarding the second question I have just mentioned. But before presenting the specific contributions I hope to make, I shall describe the general account assumed in the background of my discussion. I call it the 'broadly Kantian end-based' account of imperfect duties. This account takes imperfect duties to be requirements to adopt and promote certain obligatory *ends*.⁶ Perfect duties, by contrast, are requirements to perform or abstain from performing certain kinds of *acts*, regardless of the ends one has. Although my use of the terms "perfect duty" and "imperfect duty" tracks the Kantian tradition, I do not intend to bring my understanding of them closer to Kant's own. Doing so would be difficult at least in part because Kant does not advance a unified account of the distinction but only makes sparing remarks in the *Groundwork* and in *The Metaphysics of Morals*. How to reconstruct his account most accurately is a matter of an ongoing debate.⁷ While not engaging in exegetical work, I do draw the distinction in a broadly Kantian way.

⁶ Baron (1995: 19–107), Hale (1991: 275), Herman (2007: 213), Korsgaard (1996: 20, 137).

⁷ See, e.g., Cummiskey (1996), Denis (2001), Gregor (1963), Herman (2021), Hill (2002), O'Neill (2013), Stohr (2011).

I shall now say more about how I understand imperfect duties. I intend to outline the central features of a Kantian end-based account without taking a stance on some contentious details. This way, the arguments I advance throughout the dissertation are applicable to different kinds of end-based accounts, rather than only to some highly specific version.⁸

Obligatory ends are not to be understood as tasks or projects to be completed or discharged. Let's illustrate with an example. Consider the end of others' happiness, a paradigmatic example of an obligatory end (MS 6:385).⁹ To say that other people's happiness is an obligatory end in the relevant sense is to say that it is a "major, serious, always potentially relevant moral consideration" (Hill 2002: 206). So, someone who has others' happiness as an end is never indifferent to how her acts bear on others' happiness. Someone who has this end is committed to treating others' happiness as a source of normative reasons throughout their life. Importantly, having an obligatory end does not imply acting to promote this end on any particular opportunity. Sometimes one may act to further one's other ends (obligatory or not). What having an obligatory end does imply is having adopted or formed certain attitudes. In the case of the end of others' happiness, the relevant attitudes are the resolution to take others' happiness seriously as a source of moral reasons and the intention to promote it. This intention implies forming some more or less specific plans or policies regarding promoting others happiness' in various circumstances. Having an obligatory end also implies continuously working on cultivating some attitudes. One example is the attitude of special attentiveness to facts that are relevant to the end one has adopted. After all, if I have adopted the end of other people's happiness, I should take steps to become aware of opportunities to help others (Korsgaard 1996: 180). There are other

⁸ My presentation here is significantly indebted to Barbara Herman's (2021: Ch. 7) and Douglas Portmore's (forthcoming) discussions.

⁹ When referencing Kant's works throughout the text of the dissertation, I use "MS" to refer to *The Metaphysics of Morals* and "G" to refer to *Groundwork of the Metaphysics of Morals*. The number before the colon indicates the number of the volume, and the number after the colon indicates the page number(s) in the Berlin Academy edition.

attitudes and commitments one could reasonably be expected to work on cultivating insofar as one adopts the obligatory end of others' happiness, but the sketch offered so far should suffice for the present purposes. Importantly, having the required attitudes indirectly requires one to act to promote others' happiness. However, it is at least possible that one may have no opportunities to do so, in which case one would comply with the requirements of the duty without actually promoting anyone's happiness.

So far, I have presented the following broad picture. Imperfect duties are requirements to adopt and promote broad obligatory ends, where adoption and promotion amounts to a direct requirement on attitudes and an indirect requirement on actions. There is an immediate problem: it seems plausible that in some cases we are directly required to benefit others, not in virtue of having the relevant obligatory end. For instance, in the infamous *Drowning Child* case, where an agent has an opportunity to save a child at an incomparably low cost to oneself, there seems to be a direct requirement on action. One obvious solution is to distinguish a perfect duty of rescue from the imperfect duty of beneficence. But it is not easy to identify principled grounds for drawing the distinction (Hanser 2014: 320–321). I believe that end-based accounts have the resources to avoid this problem without resorting to positing a separate duty of rescue. One option is to claim that *Drowning Child* and analogous cases present 'golden opportunities' to help in the sense that failing to act in such cases implies a serious failure in one's commitment to the obligatory end (Noggle 2009; Portmore *forthcoming*).¹⁰ So, while agents may be strictly required to benefit others in some cases, the requirement is still indirect: it is just that a failure to act could imply a failure to form the relevant attitudes.¹¹ It can be claimed, for instance, that certain acts are necessary means to obligatory ends and thus are strictly, albeit indirectly,

¹⁰ Discussion by Sticker and van Ackeren (2018: 416) is also relevant here.

¹¹ And it is plausible that not all failures to act in cases like the *Drowning Child* would count as golden opportunities. If an agent frequently finds herself in such situations, she might be permitted to fail to benefit others in such cases due to the accumulation of costs to herself. See Pummer (2023: Ch. 6) and Timmerman (2015) for discussions.

required.¹² While I find some form of the golden opportunity condition plausible, I remain open on how exactly, if at all, it should be construed.

Although the view I have just presented is shared by most end-based accounts, many of them diverge on several important issues. The two main reasons for discussing these issues are as follows. First, this discussion shows that there are nuances that must be taken into consideration when developing a comprehensive end-based account. Second, my discussion shows that the broad account which I assume in the background of my discussion in this dissertation does not take a strong stance on any of these issues and remains relatively non-committal.

One is the question about the general relationship between obligatory ends and imperfect duties. Some accounts appear to claim that the relationship is direct, that is, each imperfect duty prescribes a specific obligatory end to adopt and promote (Hale 1991: 275; Portmore *forthcoming*; Segal 2022). These accounts can be interpreted as claiming that for each obligatory end there is one corresponding imperfect duty. Other accounts take the relationship to be more complex. They can be taken to suggest that there are more imperfect duties than there are obligatory ends (Herman 2021). On such a picture, while there is a broad, overarching requirement to adopt obligatory ends, there are many separate imperfect duties. These are the duties to cultivate different kinds of attitudes constitutive of the obligatory ends or to promote these ends in various ways. For instance, the duty to cultivate a sympathetic attitude to others can be separate from the duty to form an intention to promote others' happiness and to act on this intention.

Another issue is the question about what obligatory ends there are. Kant's own view is that there are only two obligatory ends: others' happiness and one's own perfection (MS 6: 385). Many accounts try to preserve this feature and resist positing any further obligatory ends (Denis 2001: Ch. 2, 5; Herman 2021: Ch. 7; Hill 1971, 2002: Ch. 7). These accounts also tend to take the view that the relationship between the obligatory ends and imperfect duties is complex, and that more than one imperfect

¹² See, however, an interesting argument against this proposal (Segal 2022: Ch. 2).

duty corresponds to each of the obligatory ends. For instance, imperfect duties such as those of beneficence, gratitude, and sympathy, those that Kant labels the “duties of love” (MS 6: 452), correspond to the end of others’ happiness. Other accounts do not follow Kant on the question of the number of obligatory ends. They may posit other obligatory ends. For instance, Douglas Portmore (forthcoming) discusses the obligatory end of adequately expressing gratitude to those to whom it is owed. Accounts like the one endorsed by Portmore, tend towards the view that the relationship between the obligatory ends and imperfect duties is direct, and that for each obligatory end there is one imperfect duty to adopt and promote it.¹³

My background account does not make a strong commitment to any specific views on either of the issues. I assume that there are at least two obligatory ends: others’ happiness and one’s own perfection. But it could be that there are other ends. Since throughout the dissertation I mainly focus on the obligatory end of others’ happiness, this is not a significant issue. I also do not take a stance on whether the relationship between obligatory ends is direct or whether it is more complex. Throughout the dissertation, for simplicity, I tend to discuss it as if it were direct. But my arguments would work just as well if it were presented as a more complex relationship.

There are more questions which would need to be discussed to arrive at a fully fleshed out account of end-based imperfect duties, and on which my view remains neutral. What are the grounds for obligatory ends (that is, why is any end obligatory)? How should one weigh up obligatory ends against one’s other ends? I mostly set these important substantive questions aside, only occasionally touching upon some of them at various points throughout the dissertation when they become relevant.

While I have tried to offer a sketch that should be acceptable by most of those who find some end-based view of imperfect duties attractive, I do not claim that it is a

¹³ There are also Kantian theorists who are open to the thought that there could be more obligatory ends than Kant himself admits (Baron 2016: 351).

perfectly general picture. However, the key features of the view I assume in the background of my discussion are widely accepted.¹⁴

1.3 Kantian End-Based Views in the Context of Current Debates

I shall now situate the broadly Kantian end-based view in the context of the current debates, in part to demonstrate that engaging with it is worthwhile, and in part to provide an outline of the chapters of the dissertation and explain how they contribute to the debates.

1.3.1 Imperfect Duties and Supererogation

One prominent debate where imperfect duties play an important role concerns the place of supererogation in moral theorising. This debate is specific to Kantian views on imperfect duties. It is often argued that Kantian ethics has no conceptual space for the category of supererogation at least in part due to the ‘limitless’ nature of Kantian imperfect duties. Let’s call this claim the ‘Incompatibility Thesis’.

Chapter 2 engages with this debate. I distinguish several features that could constitute ‘limitlessness’ and be responsible for the incompatibility between Kantian imperfect duties and supererogation. I then argue that none of these features generates the incompatibility. My argument relies on the distinction between ‘weak’ and ‘strong’ kinds of supererogation. One of the key problems with the arguments for the Incompatibility Thesis is that they tend to ignore this distinction. I argue that on almost any plausible view, Kantian imperfect duties are compatible with at least the weak kind of supererogation. I also argue that compatibility with weak supererogation is significant because weak supererogation does most of the work that we expect the

¹⁴ Among philosophers who understand Kantian imperfect duties broadly like this are Baron (1995), Hale (1991: 275), Hill (1971, 2002), Ignieski (2006), Noggle (2009), Portmore (forthcoming), Stohr (2011).

concept of supererogation to do. And the views that preserve the incompatibility with both weak and strong kinds of supererogation have the unsavoury implication that morality is extremely demanding.

Thus, the case for the incompatibility between Kantian imperfect duties and supererogation is much weaker than is generally thought.

1.3.2 *Imperfect Duties and Latitude*

It is widely believed that imperfect duties permit extensive latitude regarding how to fulfil them. There is, however, a debate about the kind of latitude characteristic of imperfect duties. Those who adopt end-based views on imperfect duties may also disagree about the specific kinds of latitude these duties permit and the extent to which they permit it, but they have a clear explanation of why these duties permit latitude. Since the obligatory ends are in some ways underspecified, indeterminate, or vague, they rarely prescribe specific courses of action. Agents subject to imperfect duties thus have wide latitude or discretion regarding how to act to adopt and promote these ends.

A further important question concerns the implications of accepting that imperfect duties permit latitude of some kind. It is a common view that extensive latitude permitted by imperfect duties somehow ‘limits’ the demandingness of morality (if imperfect duties are part of the moral landscape). However, how exactly the relationship between latitude and demandingness should be construed remains underdiscussed.

In Chapter 3, I aim to fill this gap. I discuss the different ways in which the relationship can be understood and argue that it cannot be plausibly understood in any interesting sense of ‘limiting’. The only plausible interpretation of the relation is as one of compatibility. That is, if imperfect duties with extensive latitude are part of morality, then it is *possible* that morality is moderately demanding (or at least not as demanding as it would be without the latitude of the relevant kind). But accepting imperfect duties as part of the moral landscape does not ‘limit’ demandingness in any of the more interesting senses. To arrive at this result, I consider and reject two potentially interesting interpretations. On the one interpretation, extensive latitude

explains limits on demandingness. On the other, latitude *implies* limits on demandingness. I argue that the former interpretation is either false or question-begging, and the latter is either false or trivially true. Given these problems, these interpretations should be rejected or set aside as uninteresting.

My discussion in Chapter 3 makes a twofold contribution to the broader debate. First, the discussion clarifies the relation between the latitude permitted by imperfect duties and demandingness of morality, which has not been satisfactorily done before. Second, the result that the relation between latitude and demandingness is of little theoretical interest is significant. On the basis of this result, we can now conclude that questions about how much latitude imperfect duties permit and how demanding morality is do not share an interesting connection. Admitting imperfect duties that permit extensive latitude in our moral theorising may help us construct a framework in which a moderately demanding version of morality could be defended. But the mere presence of imperfect duties and their latitude does not do anything to lend support to any substantive claims about demandingness.

The contributions that I aim to make to the debate on imperfect duties and latitude formally resemble the ones made in Chapter 2 where I discuss imperfect duty and supererogation. That is, in both these chapters I clarify a certain feature of imperfect duties and show that a widespread way of thinking about these duties is confused.

1.3.3 Imperfect Duties and Collective Impact Cases

A new debate involving imperfect duties has emerged in the past few decades. This debate centres on whether imperfect duties could be useful for addressing the ‘inefficacy problem’ that arises in collective impact cases. Very briefly, collective impact cases are such that when enough agents act in certain ways, a harmful or beneficial outcome occurs, but no individual agent’s actions seem to make a difference to this outcome. The inefficacy problem is the problem of identifying what it is that an individual agent is doing wrong when she participates in a collectively harmful activity, given that her actions do not make a difference to the outcome.

Recently, several theorists have proposed that agents act wrongly in collective impact cases in virtue of violating an imperfect duty (Albartzart 2019; Baatz 2014; Nefsky 2021).

In Chapter 4, I discuss whether a Kantian end-based view could form a basis for a solution to the inefficacy problem. One reason to think that it could is that it seems to satisfy the desiderata for an acceptable solution. Furthermore, a solution of this kind has already been proposed (Albartzart 2019). My aim in this chapter is twofold. First, I provide general reasons to think that a solution based on Kantian accounts would fail. Second, I put forward further specific objections to Albartzart's solution. I argue that Kantian imperfect duties do not fare well as bases for the solution to the inefficacy problem because it is possible for agents to act in collectively harmful ways, thereby contributing to harmful outcomes, without violating their imperfect duties. There are two reasons for this. One is that we can't make judgments about whether agents comply with their imperfect duties when collective harm cases are considered in isolation from how agents live their lives overall. The other is that if these cases are not considered in isolation, agents rarely if ever violate imperfect duties in virtue of their participation in collective harm cases. So, a solution based on imperfect duties would fail to condemn many intuitively wrong collectively harmful acts. I also suggest that this result is not to be regretted, as it does not by itself constitute an objection to Kantian imperfect duties.

I also consider Albartzart's version of a solution based on Kantian imperfect duties. Her key claim is that acting in collectively harmful ways fails a certain universalisation test specific to violations of imperfect duties. My strategy in arguing against this solution is to argue that depending on how specifically Albartzart's proposal is understood, collectively harmful acts would rarely if ever fail the proposed universalisation test. This implies that Albartzart's solution also yields sceptical verdicts in many cases: it fails to classify acts that collectively bring about harmful outcomes as wrong.

Chapter 5 continues the theme of collective impact cases but takes a step away from the problems concerning imperfect duties. Here, I turn to the much-discussed

question about the possibility of cases in which no individual instance of a collectively harmful act could make a morally significant difference to the outcome in question. My aim is to offer a novel argument against the possibility of such cases, at least their phenomenal variety (where the harm in question is purely experiential). Such cases have been thought to present particularly pressing challenges to those who deny their existence. If my argument is successful, then the inefficacy problem is, at least to some extent, less troubling. And this result would make it less regrettable that Kantian imperfect duties fail to provide a basis for a solution to the problem.

1.4 Concluding Remarks

In the previous section, I have outlined three broad areas of debate in which imperfect duties are relevant, and explained how my dissertation contributes to each debate. However, I do not want to claim that these debates exhaust the discussion of imperfect duties. There are many other debates — spanning from political philosophy and issues of human welfare rights to business ethics, to other issues in applied ethics — where imperfect duties are invoked.¹⁵ A thorough investigation of all the various debates would be worthwhile. It could show just how widespread imperfect duties are and how discussions of many different problems in moral theorising could benefit from a better understanding of these duties. Such an investigation, however, would go beyond the scope of this dissertation.

The fact that the concept of imperfect duty grows more and more ubiquitous is both a reason to pay serious attention to it and to be wary of the possible misconceptions and underdiscussed assumptions that might sneak into the debates. These potential problems could make the philosophical work that investigates this promising concept or that applies it to specific problems significantly less fruitful or

¹⁵ For discussions of imperfect duties and welfare rights see Hope (2014) and O'Neill (1996). Lea (2004) and Ohreen & Petry (2012) invoke the concept of imperfect duties in discussions in business ethics. There are also quite a few discussions of imperfect duties in various debates in applied ethics. One recent example is Amy Berg's paper (forthcoming) that argues for an imperfect duty to read the news.

important. It is here that my dissertation contributes. It makes progress towards dissolving some of the problems and clarifying some of the misconceptions.

Chapter 2

Kantian Imperfect Duties and Supererogation: A Case for Compatibility

2.1 Introduction

There is widespread agreement that the category of supererogation plays an important role in our moral theorising: it helps us explain moral phenomena such as heroic acts or small favours, which seem good to do, but not required by any duty. However, it has often been suggested that moral theories that include Kantian imperfect duties can and should do without supererogation.¹⁶ It is argued that these theories have no conceptual space for the supererogatory at least in part due to the nature of imperfect duties. Consider this claim by Marcia Baron (1995: 42) as an example: “The central difference between [the imperfect duty approach] and that taken by supererogationists lies not in the deeds required, but in the fact that [Kantian] imperfect duties are primarily duties to have certain maxims”.^{17, 18} And more generally, the position that imperfect duties and supererogation are incompatible is quite popular among Kantian theorists (Atterton 2007; Baron 1995, 2015, 2016;

¹⁶ Throughout this chapter, I generally drop the adjective “Kantian” before “imperfect duty”. But unless otherwise stated, my discussion covers only Kantian imperfect duties.

¹⁷ Baron seems to use the expression “duty to adopt an end” and “duty to adopt a maxim” interchangeably. The idea seems to be that to adopt a general obligatory end, you must also adopt some appropriate maxim. For instance, in the case of the end of other people’s happiness, one such maxim is that of promoting others’ happiness.

¹⁸ I shall add to this that the apparent incompatibility is not a consequence merely of the fact that imperfect duties prescribe *some* obligatory ends, but also of the fact that they prescribe certain *specific* ends such as the ends of others’ happiness and one’s own moral perfection.

Cummiskey 1996; Hale 1991; Herman 2012: 393, 2021: 161; Timmermann 2005), although there are exceptions (Hill 1971, 2002: Ch. 7; McCarty 1989).¹⁹

Specifically, it is thought that *limitlessness* of imperfect duties plays an important role in rendering them incompatible with supererogation. I shall say more about what limitlessness is and discuss it in detail in sections 2.3 and 2.4.

The aim of this chapter is to argue against the view that imperfect duties and supererogation are incompatible. Throughout the chapter, I shall call this view the “Incompatibility Thesis”. My main claim is that on any understanding of limitlessness, limitless imperfect duties are compatible with some kind of supererogation. That is, there could be a moral theory that includes limitless imperfect duties and has conceptual space for some kind of supererogation.²⁰ Establishing this claim would constitute a twofold contribution to the literature. It would make clearer the nature of imperfect duties and show that a widespread way of thinking about their relationship with supererogation is confused.

Here is how my discussion is structured. In section 2.2, I say a bit more about the Kantian view on imperfect duties and the concept of supererogation. In section 2.3, I discuss the concept of limitlessness and the apparent reasons for accepting the Incompatibility Thesis. In section 2.4, I argue that there are no good grounds for accepting the Incompatibility Thesis. In section 2.5, I offer some reflections on the significance of this result. In section 2.6, I consider several objections.

¹⁹ Since Baron (1995, 2015, 2016) offers the most detailed and explicit discussion in defence of the incompatibility between imperfect duties and supererogation, my chapter engages with her views quite extensively.

²⁰ For the sake of brevity, I often say “imperfect duties are compatible with supererogation”. Phrases like this should be read as “there could be a moral theory that includes imperfect duties and has the conceptual space for supererogation”.

2.2 Imperfect Duties and Supererogation: The Relevant Concepts

Before explaining the motivation for thinking that imperfect duties and supererogation are incompatible, I shall briefly present each concept.

2.2.1 *Kantian Imperfect Duties*

This presentation serves to help the reader recall the fuller discussion that I offered in Chapter 1. Some version of this broad account of imperfect duties is adopted by those who endorse the Incompatibility Thesis. On the Kantian account, an imperfect duty is a requirement to adopt a certain obligatory end, whereas a perfect duty is a requirement to act or abstain from acting in certain ways. Primarily, the duty prescribes adopting and promoting attitudes that constitute the relevant ends. Kant distinguishes two such ends: other people's happiness or welfare and one's own moral and non-moral perfection (MS 6:385). For instance, in the case of the end of others' happiness, the relevant attitudes would be a general intention to promote others' happiness and attentiveness to opportunities to promote it. Since this intention would be among the relevant attitudes, one would be indirectly required to form some policies regarding promoting others' happiness and to act on these policies at least sometimes. Some failures to act would be acceptable, since not every failure to act implies a failure to adopt the relevant attitudes. Some acts, however, can be obligatory. These are so-called 'golden opportunities' to act on the obligatory ends (Noggle 2009). For example, if one decides not to save the child in the standard *Drowning Child* scenario, it seems obvious that one does not have others' happiness or welfare as an end, as one fails to act on a golden opportunity.²¹

²¹ This account may diverge from Kant's own account, but that is not my concern here. What is important is that something like this account is discussed in the debates surrounding the Incompatibility Thesis.

I mainly focus on the paradigm imperfect duty of beneficence that requires helping others as part of one's commitment to the end of others' happiness. However, my discussion will also touch on the duty of moral self-improvement.

2.2.2 *Supererogation*

The standard view on supererogation that is accepted by most theorists, including the proponents of the Incompatibility Thesis, attributes at least two necessary features to the concept. One is that an act is supererogatory only if it is morally optional or, in other words, beyond duty. Put differently, an act is morally optional iff it is both permissible to do it and to abstain from doing it. This first feature is not sufficient for accounting for supererogation since some morally optional acts are morally indifferent. Choosing to sprinkle pepper on my salad is morally optional but clearly not supererogatory. The second necessary feature is that an act is supererogatory only if it is morally good, and morally better than some permissible alternative act available to the agent at the time. So, not only must supererogatory acts be morally optional, but they must also be morally good, and more so than at least one permissible alternative.²²

I do not claim that the two necessary features I have outlined are sufficient. Further conditions could be added, such as that supererogatory actions must be appropriately motivated. For instance, doing something morally excellent while being motivated solely by a desire to become famous isn't supererogatory on some views (Crisp 2013). As these further conditions would unnecessarily complicate the discussion, I shall not make any assumptions about them.

Furthermore, depending on how exactly moral optionality is understood, two kinds of supererogation can be distinguished. They are conventionally referred to in the literature as "strong" and "weak" or "unqualified" and "qualified" supererogation (Dancy 1993: 130-131; Heyd 2019; Horgan & Timmons 2010: 32–33). An act is

²² These conditions are widely accepted and uncontroversial. See, for instance, Archer (2018), Crisp (2013), Dorsey (2013), Heyd (2019), Horgan & Timmons (2010).

supererogatory in the strong sense only if it is *completely* optional, that is, it is not even pro tanto required: no moral justification is necessary for permissibly failing to perform such acts. An act is supererogatory in the weak sense only if it is optional, but *not completely* so: it is pro tanto required, but failing to perform it is justified given the circumstances.

Some people (Dancy 1993; Horgan & Timmons: 32–33) argue that only strong supererogation is the genuine one. They sometimes use the label of “quasi-supererogation” for the weak kind. Others (Rawls 1999: 100), on the contrary, endorse some version of the weak kind as *the* account of supererogation. Yet others (Portmore 2011: 131–137) recognise both kinds, claiming that there are different non-exclusive explanations of why different kinds of act might be supererogatory.

For the purposes of present discussion, I do not need to side with any of the three views. The purpose of introducing the distinction between strong and weak supererogation is as follows. As shall become clear from further discussion, all interpretations of limitlessness attributed to imperfect duties are compatible with *weak* supererogation. However, not all of them are compatible with *strong* supererogation. This, as I shall show, does not reduce the significance of my argument. The proponents of the Incompatibility Thesis often fail to appreciate the distinction between the two kinds of supererogation and make a sweeping claim that limitless imperfect duties are incompatible with supererogation *simpliciter*.²³

²³ Marcia Baron (2016: 348–349) is clearly aware of the distinction between the two kinds of supererogation since she explicitly mentions a version of the distinction. However, she claims that neither of them is compatible with the Kantian view. Moreover, she does not discuss the distinction when she advances her main arguments for the Incompatibility Thesis.

2.3 Why Think That Imperfect Duties and Supererogation Are Incompatible?

As I briefly mentioned in the introduction, the main apparent reason for thinking that imperfect duties and supererogation are incompatible is the limitlessness of the former. What exactly is limitlessness and why does it seem to render the two concepts incompatible? I use “limitlessness” as a placeholder term to refer to one or more features in virtue of which imperfect duties seem to leave no space for supererogation. While the term is not generally accepted in the literature, the features to which it refers are often implicitly attributed to imperfect duties. I distinguish three features that are thought to generate the incompatibility: *persistence*, *open-endedness*, and *requiredness*. These features are not usually explicitly distinguished, and the terminology is my own.

One feature of the relevant sort is that imperfect duties are *persistent* in the sense of retaining normative force regardless of how much one does to fulfil them. It seems that regardless of how much one does to uphold the commitment to the end of others’ happiness, one could never correctly judge oneself to be “over-and-done” with this end.²⁴ That is because such ends are not projects to be executed, but rather sets of attitudes to be maintained. In line with this, Douglas Portmore (forthcoming) suggests that an imperfect duty “doesn’t seem to be the kind of duty that you could dispatch”.²⁵ Marcia Baron (1995: 42) argues that one cannot “both have the maxim of beneficence and regard one’s “work” as over”.²⁶ The incompatibility arises because it is unclear

²⁴ See Noggle (2009) for a related discussion.

²⁵ I should note that Portmore’s view is that imperfect duties are not only compatible with supererogation, but also should leave space for this category. His argument is different from mine, however (see fn. 45 and section 2.6.2).

²⁶ The idea is that if you adopt the general end of others’ happiness, you also must adopt appropriate maxims such as the maxim of beneficence, that is, of promoting others’ happiness.

whether one could do more than is required by a duty that could not be dispatched or fulfilled.

The second feature is that imperfect duties are *open-ended* and have indeterminate conditions of compliance. This feature is a result of the obligatory ends themselves being open-ended with respect to what counts as their successful adoption and promotion. Consider a non-moral analogy. The end of leading a healthy lifestyle does not require performance of a precisely specifiable set of actions for its success. One is free to decide whether to exercise more or less and free to select among a variety of diets and sleep schedules while still staying healthy. Similarly, one may have discretion with respect to how, when, and to what extent to help others given that the end of others' happiness does not impose precise success conditions.²⁷ This feature is also generally accepted. For instance, Warner Wick (1983: li) claims that imperfect duties are "indefinite" since "what they demand is ... without assignable limits".²⁸ Given that the requirements are open-ended, supererogation might not be possible because it seems to require a clear division between what is required and what is beyond duty, and such a division is not provided by duties whose requirements are indeterminate.

Finally, there is yet a further feature that supports the Incompatibility Thesis. For a lack of a better label, this feature can be called *requiredness*.²⁹ The feature is that imperfect duties subsume all morally good actions unless these actions are already subsumed by some perfect duty. In other words, for any morally good action, there is always at least some morally requiring reason to perform it that stems from a more general requirement to adopt and promote some obligatory end, unless that reason is generated elsewhere.³⁰ Here is how Baron (2015: 229) describes it: "Saintly and heroic

²⁷ See Portmore (forthcoming) and Tenenbaum & Raffman (2012) for related discussions.

²⁸ See also Hill (2002: Ch. 7), McCarty (1989).

²⁹ The word "demandingness" would sound less clanky, but given that it is widely used in a different context and with a different meaning, I have decided against using it here.

³⁰ My discussion relies on a widely accepted distinction between *requiring* and *justifying* (or *permitting*) reasons. For a discussion of this distinction see Gert (2016) and Pummer (2023: 24–26).

actions will typically be among those actions that fall under a principle of wide imperfect duty. They are not severally required, but they are not beyond duty. They are of a type — actions promoting others' happiness — that are, as a class, morally required". If imperfect duties are limitless in the sense of having all morally good actions falling under them (apart from those that fall under perfect duties), then by definition they seem to lack conceptual space for actions that are good but that go beyond their call. And as is the case with other two features, requiredness is widely accepted among those who uphold the incompatibility (Cummiskey 1996: 109–114; Timmermann 2005: 15–23).

Why think that requiredness is a feature of imperfect duties? Baron (1995: 41–42, 97–100) offers a conditional justification. If we accept that moral perfection is an obligatory end and striving towards moral perfection is a duty, then no morally good act could be beyond duty or lacking morally requiring reasons.³¹ This is due to the nature of the duty to improve oneself morally. In a Kantian framework, adopting one's moral perfection as an end involves at least two kinds of requirement. One part of what we are required to do to fulfil this duty is to engage in introspection and self-scrutiny to see as clearly as possible the reasons that we act on and what motivates us. The other part is to actively try to cultivate the strength of our characters, trying to make moral considerations sufficient to motivate us to act, and to overcome our inclinations and

Here is a brief sketch of the distinction. If there is a requiring reason to do X, then there is a consideration that makes X at least to some extent required. This consideration may be insufficient to make X all-things-considered required. If there is a justifying or permitting reason to do Y, then it may override a requiring reason to do something else. However, having only a justifying or permitting reason to do Y does not issue any requirements to do Y in the absence of a requiring reason. Acting in accordance with morally justifying or permitting reasons is morally optional while acting in accordance with morally requiring reasons is generally not (the latter would depend on the overall balance of reasons). I use expressions "requiring reason" and "reason with requiring force" interchangeably (similarly with "justifying reason" and "reason with justifying force").

³¹ This relates to the point made in fn. 18. Imperfect duties are thought to be incompatible with supererogation not only due to their prescribing ends as opposed to actions, but also due to which specific ends are prescribed.

aversions that motivate us to act contrary to duty.³² Furthermore, if one takes the duty of self-perfection seriously, one also continuously strives to do more for others. Consider the example Baron (1995: 100) offers: “...if I am easily annoyed or depressed by unhappy people, it would be ‘good enough’ if I sent cheery greeting cards or little presents but avoided seeing or talking with them”. Insofar as we both adopt others’ well-being and our own moral perfection as ends, through self-scrutiny we would notice that our lack of empathy reveals that we are not perfectly committed to promoting others’ well-being and given that we ought to strive for moral perfection we would try to improve ourselves to correct this fault in ourselves. For example, if we take seriously our moral perfection as an end we’ll try to engage with our depressed close ones in a way that expresses our benevolence towards them beyond sending “cheery greeting cards”. The point about the connection between the end of self-perfection and others’ well-being is general: the more we strive to perfect ourselves, the more we will try to promote others’ well-being. The following quote is evidence that this is how Baron (1995: 42) understands the connection:

Hill says that according to Kant, ‘when one has satisfied [the] minimum and rather indefinite requirements [of beneficence], one may promote [others’] happiness or not, as one pleases.’ But it is not entirely as one pleases; if it were, virtue would be optional and the duty to improve oneself would be incumbent only on those who, morally speaking, are especially derelict.

The connection between moral self-perfection and beneficence is the key element in Baron’s claim that Kantian imperfect duties have the feature of requiredness and are, in virtue of this feature, incompatible with supererogation. Following Kant, Baron assumes that attaining moral perfection is impossible and therefore it is only one’s duty to continually strive towards it. Since one is under a never-ending, limitless requirement to strive to improve oneself morally, what one is required to do in terms of helping others increases as one gradually improves oneself. Attempts to impose a

³² Baron (1995: 177–178), Biss (2015), Timmons (2021: 99–101, 196–198), (MS 6:446–448).

limit on what is required rest on a misunderstanding of the nature of the duty to perfect oneself morally and Kantian ethics generally, Baron claims. So, if we have the duty to improve ourselves morally, all morally good actions are in principle morally required. Thus, accepting the Kantian duty of moral self-improvement is the ground for the feature of requiredness, which seems to render imperfect duties generally at odds with supererogation.

So far, I have outlined the broad accounts of imperfect duty and supererogation, which are accepted by those who think the two concepts incompatible. I have also presented the reasons for believing in the incompatibility: imperfect duties have one or more of the three features that seem to render them incompatible with supererogation. In the next section, I shall present my argument for the compatibility of imperfect duties and supererogation.

2.4 Limitlessness and Supererogation: Towards Compatibility

I argue that the two features I have described first, namely, *persistence* and *open-endedness*, are compatible with supererogation on both the strong and the weak views. Moreover, even if imperfect duties are *both* persistent and open-ended, they are still compatible with both kinds of supererogation. After that, I turn to discuss *requiredness*, and argue that while it is incompatible with strong supererogation, it is compatible with weak supererogation.

2.4.1 Persistence

Consider a more precise statement of persistence:

Limitlessness as Persistence (LAP). A duty's requirements are limitless only if they are persistent.³³ And they are persistent just in case they remain normatively binding even if the agent does what the duty requires of her at a given time.

In the case of beneficence, no matter how much one does to help others, it is not the case that one could ever be free from the requiring force of this duty. On a Kantian account, one is never permitted to abandon the end of others' welfare, that is, to abandon the relevant attitudes such as the resolution to promote others' welfare or attentiveness to others' needs. This is unlike the perfect duty to keep one's promises, where once a promise is fulfilled, it ceases to be a source of moral reasons. So, imperfect duties may be limitless in the sense that there is no point at which they cease to have normative requiring force.

But even if they have this feature, I argue that they are still compatible with supererogation. Two things need to be established to show that persistence is compatible with supererogation. First, it must be possible for an agent to do something that goes beyond what morality minimally requires of her where some of her duties are persistent imperfect duties. So, it must be possible to conceive of cases in which an agent has at least one morally optional act even when the agent has done all (at least for the time being) that morality requires of her. Second, doing at least one of these optional acts must be more morally valuable than doing all that morality requires of her. Since the target of my critique at this point is the view that moral theory that includes persistent imperfect duties has no conceptual space for supererogation, showing that it would have the requisite conceptual space is all I need. In other words, I only need to establish the possibility of a moral theory that includes persistent imperfect duties and supererogation.

³³ I state this as a necessary condition to leave it open whether other features need to be in place for an imperfect duty to be considered limitless in the relevant sense. Some, like Baron, seem to think that limitlessness is constituted by all three features.

So, does the presence of persistent imperfect duties rule out the possibility of morally optional acts? No, it does not. If imperfect duties are limitless only in the LAP sense, it could still be possible to have morally optional acts even when one does all that these duties require. It might be that after an agent has formed the relevant attitudes and thus has adopted the end prescribed by an imperfect duty, nothing more is required of her, even if the requirement to have the relevant attitudes is persistent. For instance, in the case of beneficence, the agent could adopt a policy of donating regularly to some effective charities, seeking opportunities to help others, and consistently acting on these opportunities. This might be enough to have the relevant commitments that constitute the end of others' happiness. Even if the agent was not permitted to abandon her policies if she wanted to fulfil the duty of beneficence, she would not be required to do more than she would already be doing. This does not rule out the possibility of further morally optional actions. For example, one could spend more time trying to find more opportunities to promote others' welfare on top of doing everything else that she does as part of her duty of beneficence (and that is, by stipulation, sufficient to count as adopting and promoting the end of others' happiness). Her spending more time on these opportunities would be morally optional.

Moreover, it is plausible that doing so would be more morally valuable than failing to do so. It could be that no alternative available to her is better supported by moral reasons. So, it could be that the morally best option for her would be to form the attitudes her imperfect duty directs her to form, and on top of that to continue looking for opportunities to promote others' welfare, even though doing so would not be necessary for complying with the duty. This does not make the duty any less persistent: at no point is she permitted to stop doing whatever is necessary for upkeeping the relevant attitudes and commitments that are constitutive of the obligatory end.

To drive the point home, consider a further point that LAP is not unique to imperfect duties. Requirements of perfect duties not to kill and not to harm are persistent in the same way. No matter how successfully one avoids harming others throughout one's life, there never comes a point at which this duty ceases to have

requiring force. Importantly, while these duties have LAP, one could still be doing all that they require: once an agent abstains from harming or killing one, she is not required to do anything else even if the requirements are persistent. So, one is not required to do anything beyond abstaining from harming others even if the requirement not to harm never loses its normative force.³⁴ But it would be even more of a stretch to argue that a moral theory that includes such persistent perfect duties has no space for supererogation. It is clear that many moral theories that include such ordinary conceptions of the duties against harming and killing are compatible with supererogation. Persistence of requirements, then, does not seem to be the feature that renders duties incompatible with supererogation.

I have argued that persistent imperfect duties are compatible with supererogation but have not yet said anything about the kinds of supererogation. I shall now show that duties with LAP are compatible with both strong and weak forms of supererogation.

Recall that an act is weakly supererogatory only if it is *pro tanto* required but not all-things-considered required because the agent is justified in failing to perform it. Let's consider again the example of an agent who looks for opportunities to promote others' welfare on top of doing what her duty of beneficence minimally requires. Recall also that looking for these opportunities would be morally better than any available alternative. It may very well be that there is some requiring reason that favours looking for those opportunities but that there is some justifying factor that outweighs the strength of this reason.³⁵ In such a case, while looking for those opportunities would be part of the duty without the relevant justification, the presence of this justification would render the action weakly supererogatory.

Duties that have LAP are also compatible with strong supererogation. One could coherently imagine that there is no requiring reason for doing more than what is

³⁴ While I have used negative duties in the example, this does not seem to make the point less significant. Furthermore, negative duties such as the duty not to harm may often require positive actions on one's part.

³⁵ See Portmore (2011: 131–136) for a related discussion.

minimally sufficient for adopting and promoting the obligatory ends (such as the end of others' happiness). So, it could be that looking for further opportunities to promote others' welfare beyond doing all that the agent does to adopt and promote the obligatory end is completely optional. And since we have also assumed that it is morally better than any alternative, it is then strongly supererogatory.³⁶

So far, I have argued that a moral theory that includes persistent imperfect duties has conceptual space for supererogation of both weak and strong kinds.

Before moving on to discuss the remaining two features, there is an objection to consider. One may argue that there could be no way to go beyond persistent duties, no way to overfulfil them. For instance, in the case of the (perfect) duty not to harm, as long as an agent is doing all that this duty requires, it is impossible to do more. That is, it is impossible to do something more morally valuable than merely abstaining from harm that would count as *overfulfilling this particular duty*. Something similar could be argued in the case of persistent imperfect duties. It might be impossible to do something more morally valuable than is required by the duty of beneficence in a way that would count as overfulfilling this very duty. My response is that even if this were the case, this would not make my argument unsound. It is irrelevant whether our doing something more morally valuable than what we're minimally required to do may be categorised as overfulfilling some specific duty. This is because my aim has been to argue that a moral theory that includes persistent imperfect duties is compatible with supererogation, and *not* that any given persistent duty could be overfulfilled. So, even if overfulfilling the persistent duty of beneficence was impossible, it could still be possible to do something supererogatory. It just would not count as overfulfilling this specific duty. Furthermore, we might not want to grant that overfulfilling persistent imperfect duties is impossible in the first place. There is nothing incoherent about the

³⁶ A proponent of strong supererogation would have to explain what moral reasons speak in favour of the strongly supererogatory act if these reasons have no requiring force. But this is not a concern for my argument. My aim is merely to establish the possibility of a moral theory that has both limitless imperfect duties and supererogation. Whether such a theory would be plausible is a separate discussion, which I touch upon in section 2.6.2.

idea of promoting others' well-being to a greater extent than is minimally necessary to count as having the relevant attitudes constitutive of the obligatory end.

2.4.2 Open-Endedness

Let us now turn to another feature that could constitute limitlessness, namely, open-endedness. Here is how it can be stated:

Limitlessness as Open-Endedness (LOE). Imperfect duties are limitless only if they are open-ended. Their requirements are open-ended just in case their conditions of compliance are indeterminate.

What is meant by “indeterminacy” in the context of limitless imperfect duties? Generally, it is indeterminate that q iff there is no fact of the matter whether q . In other words, it is neither true nor false that q . A duty's conditions of compliance are indeterminate just in case there are courses of action for which it is indeterminate whether taking them is sufficient for complying with the requirements. By contrast, a duty's conditions of compliance are determinate just in case all possible courses of action are neatly sorted into two categories: courses of action such that taking any one of them would be sufficient for complying with the duty and those such that taking any one of them would not be sufficient. As an example of a duty of the latter kind, consider a duty to return the £10 that I've borrowed from my friend. As an example of the former kind, consider the duty of beneficence. It might be neither true nor false whether I am doing enough to count as having adopted the end of other people's happiness.³⁷

Now, there is a clear way to argue that the limitlessness of imperfect duties (in the LOE sense) is compatible with both weak and strong kinds of supererogation.

It should first be noted that for any duty that has indeterminate conditions of compliance, there must be at least one course of action such that it is determinately

³⁷ A similar example, called *Mr. Borderline*, is discussed by Portmore (forthcoming).

true that taking it would determinately be sufficient for complying with the duty and at least one such that taking it would determinately be insufficient. This is because if there were no such courses of action, duties with indeterminate conditions of compliance would fail to be action-guiding. They would fail to be action-guiding because it would then be impossible to ever judge whether an agent complies with the duty or not. Taking an example from Tom Dougherty (2016: 450), the situation is similar to the one where we are required to sort a bunch of tiles into two piles. We are required put a tile in the left pile only if it is red and to put it in the right pile only if it is not red, but all the tiles that we are given are neither determinately red nor determinately not red. The moral requirement regarding sorting the tiles does not then give any guidance as to what we ought to do. So, it must be the case that even for duties with indeterminate conditions of compliance, there are courses of action that are determinately sufficient for complying and courses of action that are determinately insufficient.

But once we grant that there must be courses of action that are determinately sufficient, we open conceptual space for doing more than that. For instance, in the case of beneficence, suppose that forming a certain policy of helping others and sticking to this policy is determinately sufficient for complying with the duty. It is conceptually possible that an agent could do something over and above sticking to this policy and that doing so could be more morally valuable than sticking to the policy. Now, if going over and above the policy was also pro tanto required (perhaps in virtue of a general requirement to promote the good), doing so would be supererogatory in the weak sense. If going over and above was not even pro tanto required, doing so would be supererogatory in the strong sense. Importantly, both kinds of supererogation are compatible with imperfect duties being limitless in the LOE sense.

I should emphasise that I do not claim that any moral theory that has open-ended imperfect duties would *in fact* have any supererogatory acts. It might be, for instance, that going over and above the policy would never be more morally valuable than sticking to it. But, as I have mentioned previously, my goal is to establish conceptual compatibility rather than to show that any specific moral theory that has imperfect

duties limitless in one or more of the relevant senses does in fact admit of supererogation (although I find it plausible that many do).

I have argued that a moral theory with limitless imperfect duties — in either the LRP or LOE senses — is compatible with both weak and strong forms of supererogation. Would things be different if limitlessness was constituted by a conjunction of LRP and LOE? On this view, a duty would be limitless only if it was persistent and open-ended. The answer is no. Such a duty would still be compatible with both kinds of supererogation. Take again the example of my policy of helping others, that I adopt in order to comply with the duty of beneficence. Let's assume that this duty is open-ended (perhaps sticking to my policy is determinately sufficient for complying with the duty, but choosing a less effective policy would not be determinately sufficient). Let's also imagine that the duty is persistent (I am never free from the requirement to pursue my policy). But even if both these assumptions are made, both weak and strong kinds of supererogation remain possible. Doing more than sticking to my policy could be a genuine opportunity and it could be more morally valuable than sticking to the policy.

2.4.3 Requiredness

Limitlessness as requiredness yields a stronger version of the Incompatibility Thesis (as opposed to LAP and LOE). It is stronger because imperfect duties that are limitless in this sense are incompatible with the strong kind of supererogation. Nevertheless, in this section I argue that imperfect duties having this feature are compatible at least with weak supererogation. Here is how requiredness can be formulated:

Limitlessness as Requiredness (LAR). Imperfect duties are limitless just in case all morally good actions fall under them (unless they fall under some different duty).

Now, how should the claim that all morally good actions fall under imperfect duties be interpreted? One straightforward interpretation is that these actions are pro tanto required. In other words, there are moral reasons that play a requiring role and that speak in favour of performing these actions. If this interpretation is correct, then imperfect duties that are limitless in this sense are clearly incompatible with strong supererogation. That is because, for strong supererogation to be possible, *completely* optional actions that are morally better than what is minimally required must be possible. But on the view that imperfect duties have the feature of requiredness, all morally good actions are at least pro tanto required, and so no morally good action is *completely* optional. I do not see a different plausible interpretation of requiredness that does not involve claiming that all morally good acts are pro tanto required. Moreover, some of the claims that Baron makes strongly suggest that this is what she takes to be the feature responsible for the incompatibility. Consider again the quote I have mentioned above (Baron 2015: 229): “Saintly and heroic actions will typically be among those actions that fall under a principle of wide imperfect duty. They are not severally required, but they are not beyond duty. They are of a type — actions promoting others’ happiness — that are, *as a class, morally required*” (emphasis mine). Her claim that these actions are “required as a class” can be very plausibly interpreted that she thinks that they are at least pro tanto required. There are other claims to the same effect, such as that “imperfect duties cannot be claimed to be seen as admitting of a plateau, a point beyond which more conduct of the same sort is supererogatory” (Baron 1995: 42) and that “on a Kantian view, there is just no ‘beyond duty’” (Baron 2016: 356).

Even if all morally good acts are pro tanto required, imperfect duties that have LAR are compatible with weak supererogation. This is because it does not follow that failing to perform all these acts would amount to failing to comply with one’s imperfect duties. Baron (1995: 42) seems to agree with this, claiming that successfully adopting the end of others’ happiness does not require agents to maximise the good. One could comply with the requirements of beneficence (and other imperfect duties) while failing to do as much good as one can provided that one has justification for one’s failure. But

in such a case, doing more good than is required by the duty of beneficence could be considered weakly supererogatory. So, even if a moral theory includes imperfect duties that have LAR, this moral theory still has conceptual space for weak supererogation.

The following example, which I borrow from Baron (1995: 52–53) should make the compatibility clearer³⁸:

Jill and Maria. Imagine Jill and Maria, both of whom have adopted a maxim of beneficence. Both, that is, really care about the welfare of others... <...> Jill helps others often, but she does not make the sacrifices of time and energy that Maria makes. Maria volunteers at a rape crisis centre. Jill does not turn her back on needy people and is sensitive to the needs of others, but she does not go out of her way to involve herself in activities of the sort that Maria takes part in. In addition, Jill tends not to take notice of others' needs except when she is personally acquainted with the people in question, or when a crisis such as a famine on a distant continent receives much more attention in her community than faraway crises usually receive.

It is assumed that what Jill does is sufficient to comply with the duty of beneficence. After all, she adopts the relevant obligatory end. But if the duty of beneficence has LAR, then promoting the good to a greater extent is pro tanto required. But since doing as much as Jill does is sufficient for complying with the duty, it follows that there is some justification for not doing more (the exact justification is not specified). So, when Maria does more than is minimally sufficient for compliance, she does something pro tanto required, but something she is justified in failing to do.

³⁸ Baron's assessment of the case is quite different. Specifically, she claims that rather than demonstrating compatibility, this case actually shows that we do not need the concept of supererogation to explain the difference between Maria's and Jill's actions.

So, she does something supererogatory. Thus, there is a case of an imperfect duty that has LAR but that is also compatible with weak supererogation.

An objection could be raised at this point. If imperfect duties have LAR, then no morally good action could go beyond duty. Yet, by recognising that what Maria does is supererogatory, we are claiming that her actions do go beyond duty. There seems to be a contradiction in the position I am advancing.

There is no contradiction, however. This objection would be based on a failure to distinguish between weak and strong supererogation. If we did not draw the distinction, we would not be able to coherently claim that the duty of beneficence has LAR and that Maria does something supererogatory. For then we would have to claim that no morally good action goes beyond duty and that some morally good actions do. But the two forms of supererogation imply two different senses of “going beyond duty”. It is true that if imperfect duties have LAR, no morally good action could go beyond duty in the sense implied by the strong kind of supererogation. Going beyond duty in that sense amounts to an action’s not being even pro tanto required. But an action could go beyond duty in the sense implied by the weak kind of supererogation, because for an action to go beyond duty in that sense is nothing else but to be pro tanto required but not required all-things-considered. It is beyond duty in the sense that it need not be performed for successfully complying with the duty. And this is compatible with the duty having LAR, since having this feature amounts to the claim that all morally good actions are pro tanto required.

The incompatibility between imperfect duties that have LAR and weak supererogation, then, is at most terminological. One could try to deny the characterisation of Maria’s actions as in any way supererogatory and claim that they are simply actions that are generally pro tanto required by the duty of beneficence. But that would not be an objection against the compatibility. Substantively, there would be no difference between saying that what Maria does beyond what is minimally required for compliance with the duty of beneficence generally falls under her duty of beneficence and that what she does is weakly supererogatory.

At the end of section 2.3, I claimed that on Baron's view, the duty of moral self-improvement provides the reason for thinking why imperfect duties generally have LAR. I have already offered a general argument for why LAR is not a problem for compatibility with weak supererogation. So, claiming that imperfect duties have LAR because moral self-improvement is a duty does not lend support to the Incompatibility Thesis. But it is worth showing explicitly that my argument extends to the duty of moral self-improvement given the importance of the obligatory end of moral perfection in Baron's argument for the Incompatibility Thesis.

Even if we grant that we are required to continually strive to improve ourselves morally, this does not give a reason to reject the possibility of supererogation. We could also grant, in accordance with LAR, that each act that contributes to one's moral perfection is *pro tanto* required. But even if one is required to continually strive to moral perfection, one is not required to take every action that could contribute to this end. It is possible and plausible that some acts of this kind would not be all-things-considered required, despite being morally better than permissible alternatives. Importantly, this does not imply that one is at any point justified in judging oneself to have become morally perfect or at least morally good enough so that further efforts are not required. The minimal conditions of compliance with the duty of self-improvement could be such that one is required to continually exert efforts to improve oneself. But they could also be such that as long as one does so wholeheartedly and exhibits a clear pattern of improvement (what sort of pattern that should be is a separate substantive question), one is not required to do anything beyond that. This is not to say that one stops improving at a certain point. The idea is that as long as one maintains a certain level of effort to improve oneself, one complies with the requirements of the duty. This implies that one may never judge oneself to be good enough, and on the contrary, one may always judge oneself to be short of perfection and lacking in virtue. But since the requirement is only to strive to perfection, and not to reach it, one could be coherently and plausibly judged to be complying with the requirement of beneficence if one does

what is required to count as successfully striving.³⁹ Doing more, or striving more intensively, could then count as at least weakly supererogatory (as something pro tanto required and morally better than alternatives, but something that one is justified in failing to do).

Furthermore, if we accept that the duty of self-perfection leads one to try to do more to promote others' well-being, we could write this requirement into the compliance conditions of the duty of beneficence. We would say that one cannot be correctly judged to have adopted both the ends of moral perfection and of others' well-being unless one gradually increases efforts, attention, or resources that one devotes to promoting others' well-being as one develops morally.⁴⁰ And again, this does not in principle preclude the possibility of weak supererogation.

So far, I have argued that imperfect duties are compatible with supererogation regardless of how their limitlessness is understood. If it is understood as persistence or open-endedness, they are compatible with both weak and strong kinds of supererogation. If it is understood as requiredness, they are compatible at least with weak supererogation.

2.5 Further Steps: Importance of Weak Supererogation

The result that imperfect duties with LAP or LOE are compatible with both forms of supererogation seems satisfying. The Incompatibility Thesis is false when limitlessness is understood in either of these two senses. However, the result in the

³⁹ This seems continuous with what Kant writes about the duty of moral self-perfection (MS 6:446–447).

⁴⁰ Dougherty (2017) follows a similar line of thought. His justification for the duty to develop oneself morally, however, is very far from Kant's. For Dougherty, the requirement to improve oneself is grounded in the requirement to promote others' well-being, while for Kant these are independent obligatory ends that one ought to embrace. See also Flescher (2003) for a related discussion.

case of imperfect duties with LAR may seem less satisfying. Is the compatibility between imperfect duties with LAR and weak supererogation significant? Yes, it is significant.

The defence of this answer relies on arguing that much, and perhaps all, of the work that the concept of supererogation is expected to do can be done by weak supererogation. If that is true, securing the compatibility of imperfect duties with LAR with weak supererogation is a significant result. And, in such a case, failing to establish the compatibility with strong supererogation would not be as regrettable. Making this argument, of course, would require taking a side in the debate about which kind of supererogation is more important (and up to this point I have not taken sides in this debate). If we take this argumentative route, we can no longer claim that strong supererogation is more important.

One strong reason for having the concept of supererogation is that we want to distinguish actions that are especially morally good, but which are not wrong to fail to perform. A classic example of such an action that is often discussed in the literature is the decision of a doctor to travel to a plague-ridden city to help those remaining there, on the assumption that this doctor has no special obligation to do that (Feinberg 1960: 280; Urmson 1958: 201–202). It is widely accepted that the doctor's action is supererogatory. But the analysis of this case would differ depending on whether this action is strongly or weakly supererogatory. And there are reasons to believe that analysing the doctor's action as supererogatory in the weak sense is at least as plausible as analysing it as strongly supererogatory.

Let's consider each analysis in turn. If the doctor's action is strongly supererogatory, failing to travel to this city to help people there is not even *pro tanto* required. So, an advocate of strong supererogation would have to argue that there is a moral reason that favours the doctor's travelling to this city but that has no requiring force. A defender of weak supererogation would argue that there is a morally requiring reason to travel to the city. However, the strength of its requiring force is outweighed by some justification. Given the risks involved in such an action and assuming that morality does not require agents to bear unreasonably high costs, it might be simply

too costly for the doctor to do this. There, of course, could be a more interesting story about what exactly justifies failing to help the plague-stricken, but that is beside the point.

Here's why the analysis of the doctor's action as weakly supererogatory is attractive. If his action is weakly supererogatory, we preserve all the core features of the concept of supererogation. This action is optional in the sense that it would not be all-things-considered wrong to do (even though it is *pro tanto* wrong not to do). It is also safe to assume that it is the morally best action available to the doctor at the time. In section 2.2, I have suggested that there may be further conditions for an action to count as supererogatory. Although I do not assume that any such further condition is necessary, it would be good to show that weakly supererogatory actions can satisfy some of them. I have mentioned that on some views, supererogatory actions must be appropriately motivated (Crisp 2013). For instance, if the doctor travelled to the city to become famous, his action, on these views, would not be supererogatory. But it is plausible that his action is motivated by a genuine commitment to helping others rather than by any self-interested considerations. It is also often suggested that supererogatory actions are characterised by a certain asymmetry in the appropriateness of reactive attitudes. It is thought that performing supererogatory actions is praiseworthy, but failing to perform them is not blameworthy (Horgan & Timmons 2010: 32; McNamara 2011: 202–204).⁴¹ Weak supererogation can accommodate that. It is plausible that the doctor could be appropriately praised for choosing to go to the plague-ridden city even though he is justified in failing to do that. Similarly, given the justification, it is plausible that blaming him would not be appropriate had he decided not to go.

Furthermore, it could be argued that there are certain theoretical advantages in treating weak supererogation as *the* concept of supererogation. If there are good reasons to view the concept this way, then the incompatibility with strong supererogation is not to be regretted at all. One advantage of weak supererogation is

⁴¹ See, however, Archer (2016) for a critique of this claim.

that it lets us preserve the plausible view that there is a pro tanto moral requirement to save others from non-negligible harm. That is, there is always a morally requiring reason to save others from harm whenever we can do so. This reason can be weak, and in many cases it would be outweighed by other considerations (cost of saving others, for instance). But importantly, this reason, regardless of its strength, is always there. By contrast, proponents of strong supererogation would have to explain why in some cases saving others from harm is supported by requiring reasons whereas in others they are supported by some other kind of favouring, but non-requiring reason.⁴² Alternatively, they could agree with the defenders of weak supererogation that there is always at least some morally requiring reason to save others from harm. But then they would have to admit that saving others from harm is never truly supererogatory, that is, never supererogatory in the strong sense.⁴³

I have argued that the compatibility between imperfect duties with LAR and weak supererogation is significant because weak supererogation can do the work that we expect from the concept of supererogation. Furthermore, I have offered a reason why weak supererogation might even be theoretically preferable to strong supererogation.

2.6 Objections

2.6.1 LAR and Weak Supererogation

I have argued that imperfect duties that have LAR are compatible with weak supererogation and have used the *Jill and Maria* case to illustrate this point. It could be objected that the *Jill and Maria* case could be analysed differently. Some Kantian theorists (Atterton 2007; Cummiskey 1996; Timmermann 2005) argue that imperfect duties generally require one to perfect oneself and to promote others' happiness as

⁴² Horgan and Timmons (2010) provide an account of such a favouring but non-requiring reason. See also Portmore (2011: 131–137).

⁴³ There is a third alternative. They could argue that saving others from harm is never morally required, but this seems to be a clearly implausible argumentative route.

much as it is in one's power. Failing to act on one of the obligatory ends could only be justified if doing so violated some perfect duty such as the duty not to harm others. But in such a case, failure to contribute to obligatory ends would not only be justified but also obligatory. On such views, whenever one has an available act that is morally better than all the alternatives at the time, such an act would be obligatory unless prohibited by any perfect duties. Such views, then, leave no space even for weak supererogation. For, on these views, a pro tanto required act that is morally better than alternatives could never be optional: it is either required or prohibited. So, on these views, Maria's actions would be obligatory while Jill's actions would be prohibited as insufficient for complying with her duties.

I agree that on such a view, there would be incompatibility between imperfect duties with LAR and weak supererogation. But this is not a problem for my argument. The objection relies on a certain substantive assumption about what could count as a justification for failing to promote the good. The assumption is that one's failure to act on the obligatory ends could be justified only if acting on them would violate a perfect duty. But this assumption is not a necessary part of the view that includes imperfect duties that have LAR. It is not explained or derived from facts about imperfect duties or their features. So, the incompatibility arises not because of imperfect duties themselves, but because morality is very demanding — and the reasons for its demandingness are independent of facts about imperfect duties.⁴⁴ There is also no decisive case for accepting this assumption. Why should we restrict the scope of admissible justifications for failing to do what is pro tanto required? Why not keep it in a way that preserves the compatibility with weak supererogation?⁴⁵

⁴⁴ I discuss the relation between the nature of imperfect duties and demandingness of morality in Chapter 3.

⁴⁵ The line of argument in my response can be taken further. It can be questioned whether LAR itself is an essential feature of Kantian imperfect duties, and not a further substantive assumption. It is conceivable that imperfect duties are persistent, open-ended, but not all morally good acts fall under them. If we rejected LAR, then limitless imperfect duties would be compatible with both strong and weak forms of supererogation. This view on imperfect duties would also be closer to Hill's (1971) and

One reason that the proponents of the objection could put forward is that this restriction is closer to Kant's own view (after all, they are Kantian theorists). Indeed, at times Kant does seem to suggest that the scope of admissible justifications for failing to promote the obligatory ends is narrow. For instance, he (G 4:430) claims that one does not fully embrace humanity as an end in itself if one does not try "as far as [one] can advance the ends of others". There is another often cited passage, where Kant (MS 6:390) claims that even though imperfect duties permit latitude regarding how to fulfil them, this latitude does not amount to a "permission to make exceptions to the maxim of actions, but only as a permission to limit one maxim of duty by another..."

It is not obvious that the more restrictive view is a more accurate interpretation of Kant's views. There are passages in Kant's texts that seem to support a less restrictive scope. For instance, Kant (MS 6: 393) claims that while one ought to promote others' welfare, "...how far [one's sacrifices] should extend depends, in large part, on what each person's true needs are in view of his sensibilities, and it must be left to each to decide this for himself". If one is allowed to include one's true needs in determining the weight of reasons and "true needs" are understood broadly enough, the scope of justification could be wider than the proponents of the restrictive view suggest.

This is of course, far from a satisfactory discussion of Kant's view. But that is not a problem. More importantly, even if the more restrictive view is closer to Kant's view, this does not provide a reason to accept it. After all, the point of the present discussion is to arrive at a more plausible understanding of the scope of admissible justification for failing to promote the obligatory ends and not to arrive at the most accurate version of Kant's view. While using resources from Kant's ethics, my argument is not committed to his view.

There is a stronger reply to the problem of choosing between a less restrictive and more restrictive view. The more restrictive view yields a far more demanding version

Portmore's views (forthcoming). I have not pursued this line of argument because my strategy has been to argue that even if we grant that imperfect duties are limitless in the sense close to that adopted by the proponents of the Incompatibility Thesis, the Thesis is still false.

of morality, since it introduces a general requirement for promoting the obligatory ends in the best way where nothing but violation of some other duty could justify failing to promote the obligatory ends. Personal projects and concerns, if they are not part of the duty, could not count as considerations that could justify failing to act on the obligatory ends. This would seem quite demanding for many people. And this may be seen as a *prima facie* reason for preferring a less restrictive view. In the absence of a further reason to opt for a more restrictive view, we should settle for the view on which imperfect duties that have LAR are compatible with weak supererogation.

2.6.2 *Mere Compatibility?*

So far, I have argued for the conceptual compatibility of limitless imperfect duties and supererogation. But this only establishes that moral theories that include both imperfect duties and supererogation are not incoherent. But do we have any reason to prefer a theory of this kind? One could claim that while there *could* be moral theories that include both imperfect duties and supererogation, we have more reason to prefer some moral theory on which there is no compatibility. If this objection succeeds, my argument, while sound, is not very significant.

Offering a comprehensive defence of supererogation-friendly theories would take the discussion far beyond the scope of this chapter. But even without such a defence, there are strong reasons to think that theories of such kind are not mere possibilities but serious contenders to be the kind of theory we should accept.

One important reason is that these theories yield intuitive verdicts about cases. Consider the *Jill and Maria* case again. It is intuitive that Maria is justified in failing to volunteer and to do less, perhaps to do just as much as Jill does. It is also especially good that Maria engages in volunteering. We may even accept that she would be required to volunteer if it were less costly for her. All these plausible verdicts are possible if we accept a theory that has space for imperfect duties with LAR and weak supererogation. Furthermore, our assessment of appropriate reactive attitudes such as blame or praise would be intuitive in this case. We would conclude that Maria is

praiseworthy for volunteering, but would not be blameworthy if she failed to do that. After all, Jill is not blameworthy for doing less, provided that she does enough overall.

Another reason is that supererogation-friendly theories that include limitless imperfect duties are seriously considered in the literature and receive significant support. For example, Thomas Hill Jr. (1971, 2002) offers an interpretation of Kantian imperfect duties that have the features that I call persistence and open-endedness and that are compatible at least with weak supererogation. Portmore (forthcoming) argues for a similar claim, although his account of Kantian imperfect duties is somewhat distinct from Hill's. Moreover, Baron's own proposed view of imperfect duties that have the feature of requiredness is compatible with weak supererogation, as has been argued in the previous section. While I have not provided reasons to accept any of the views I have just mentioned, what I have said should be sufficient to ward off the objection that the supererogation-friendly theories are mere logical possibilities but not views worthy of serious consideration. I don't claim that one of these views must be correct but rejecting them without argument would be unwarranted.

Thus, the charge that the result that I have reached is not very significant is unconvincing: the kinds of theories that, as I have argued in this chapter, are coherent, are also plausible and not just mere logical possibilities.

2.7 Conclusion

In this chapter, I have argued that a moral theory that includes limitless imperfect duties has conceptual space for supererogation. I have first discussed the three features that are thought, jointly or separately, to be constitutive of limitlessness. This is already worthwhile, since, as far as I am aware, there has been no explicit discussion of them so far in the literature. It is important to have established what exactly serves as the reason to think that imperfect duties leave no space for supererogation. I have then shown that none of these features renders imperfect duties incompatible with supererogation. Specifically, if limitlessness is understood as persistence (LAP) or open-endedness (LOE), imperfect duties are compatible with both weak and strong

forms of supererogation. If limitlessness is understood as requiredness (LAR), imperfect duties are compatible only with weak supererogation. I have then argued that compatibility even with weak supererogation is significant, since there is a good reason to think that weak supererogation can do the relevant conceptual work.

Furthermore, while I have argued only for the compatibility between imperfect duties and supererogation, there are also reasons to think that theories that include both supererogation and imperfect duties are not mere possibilities, but should be taken seriously.

Finally, I should mention a limitation to my result. While I have argued for the compatibility between *Kantian* imperfect duties and supererogation, my argument does not establish that *Kant's ethics* can accommodate supererogation. There are other reasons, apart from those grounded in the nature of imperfect duties, that could serve as the source of the incompatibility. So even though imperfect duties do not pose an obstacle for accommodating supererogation, some other elements could.⁴⁶

⁴⁶ See, for instance, Guevara (1999) for an argument that Kant's ethics cannot accommodate supererogation because of the theory of value Kant adopts.

Chapter 3

Imperfect Duties, Latitude, and Demandingness: Searching for the Relation and Clarifying the Misconceptions

3.1 Introduction

It is widely thought that one of the key distinctive features of imperfect duties is that they allow the agent latitude in how to act in accordance with their requirements.⁴⁷ For example, in the case of the duty of beneficence, a paradigm imperfect duty, agents have latitude regarding whom to benefit, on which occasions, and perhaps even to which extent. Perfect duties, by contrast, permit far less extensive latitude. For example, take the paradigmatically perfect duty to keep one's promises. One may choose between returning money one promised to return in cash or by a bank transfer, but the timeframe, the recipient, and the amount could plausibly be precisely specified by the duty.

It seems plausible that the fact that imperfect duties permit extensive latitude is in some way connected to the facts about demandingness of morality. If I am free to choose when, how, whom to benefit, and even to what extent, then I could choose ways of helping that are more convenient for me and require me to sacrifice less for others. This line of thought is taken by quite a few philosophers. For instance, Patricia Greenspan (2010: 184) reasons in a very similar way:

⁴⁷ This is true across virtually all accounts of imperfect duties, and so my discussion in this chapter does not presuppose a specific account.

...the notion of imperfect obligation is important to a common-sense deontological approach precisely as a limitation on morality's binding demands. We need only contribute a reasonable amount of aid to the needy, say — at some vaguely specified level, possibly rather high, but still leaving us leeway to choose when and to whom and how much we contribute, and hence how much we have left over to devote to the pursuit of our optional ends.

Furthermore, she adds that “...imperfect obligation ... can limit the moral demands on an agent by leaving her leeway for choice” (Greenspan 2010: 204). Similarly, George Rainbolt (2000: 249) claims that latitude permitted by some imperfect duties “keeps the obligation ... from being too binding, too demanding”. Simon Hope (2014: 405) suggests that perfect duties “contain no discretionary element: they are often very demanding”. This seems to imply that duties that do contain such a discretionary element, or that permit latitude, would not often be as demanding. These are just some examples of the claims made in the literature.

But what exactly is the relation between latitude and demandingness? Can we specify the relation beyond the imprecise remarks that latitude “limits” demandingness or “keeps” morality from being too demanding? These questions have not yet received a satisfactory treatment. In this chapter, I aim to fill this gap. As I shall argue, the relation is much weaker and much less interesting than may appear. I shall argue that there is no theoretically interesting or significant sense, in which latitude could “limit” demandingness or “keep” morality from being very demanding. On any interpretations that seem theoretically significant, such claims are false, or question-begging, or trivially true. Thus, no interpretation that seems theoretically significant is in fact theoretically significant.

The best way to understand claims about the latitude “limiting” the demandingness is to understand them as the claim that imperfect duties that permit extensive latitude are *merely compatible with* moderately demanding morality. But whether morality with imperfect duties is overly demanding does not directly depend

on the latitude permitted by imperfect duties. Their latitude does not “limit” demandingness in any interesting sense.

My discussion is organised as follows. In section 3.2, I introduce the concepts of demandingness and latitude. In section 3.3, I distinguish two potentially theoretically significant senses, in which the relation between latitude and demandingness could be understood. In sections 3.4 and 3.5, I argue that the claims about the relation in these senses are false, or question-begging, or trivially true and thus should be rejected or set aside as insignificant. In section 3.6, I present what I take to be the correct construal of the relation. In section 3.7, I respond to a potential objection.

3.2 Demandingness and Latitude

3.2.1 *Demandingness*

Demandingness is standardly understood in terms of cost. On this view, it is understood as a function of costs or sacrifices that an agent is required to bear to fulfil a given moral duty.⁴⁸ So, the costlier it is for an agent to satisfy the duty’s requirements, the more demanding the duty is for the agent. Incurring or bearing costs makes one worse off than one otherwise would have been (unless one receives an appropriate compensation).⁴⁹ For instance, losses to one’s welfare, setbacks to one’s projects,

⁴⁸ See, for example, Kagan (1989, esp. Ch. 1 and Ch. 7) and Murphy (2000: 16–17).

⁴⁹ This resembles the standard counterfactual analysis of harm. The question of whether harm and cost are in any way distinct is an interesting one. Intuitively, there is some difference. For instance, when I buy a cake in a bakery, the money that I pay is a cost that I incur, but it is counterintuitive to say that I am harmed by paying this money. It seems that at least one difference is that talking about cost makes sense only in the context of trade-offs, whereas harm is not confined to this context. For instance, in the case of the discussions of demandingness of morality, costs are reductions in well-being one undergoes to increase others’ well-being or, perhaps more generally, to do the right thing.

frustration of one's desires, pain, and unpleasantness all count as costs.⁵⁰ Throughout the chapter, I shall only discuss demandingness in terms of cost.⁵¹

In the literature on demandingness of morality, views like Shelly Kagan's (1989) and Peter Singer's (1972) are often criticised for being very demanding. Their views can be taken to imply, for instance, that spending money on personal entertainment and cultural activities would often be, given the current state of the world, morally unjustified. Furthermore, their views may imply that we are not morally entitled to choose our occupation since some career paths are vastly more effective in promoting the good or preventing suffering. When I say that morality is extremely demanding, the level of demandingness advocated for by Singer or Kagan is broadly what I have in mind. And when I say that morality is moderately demanding, I refer to some level of demandingness that is significantly lower than the extreme one, but not counterintuitively low.

3.2.2 *Latitude*

I now turn to clarify the concept of latitude. Generally, a duty permits latitude just in case there are at least two distinct courses of action such that taking any of these courses of action is sufficient to satisfy the requirements of the duty. This account extends to both perfect and imperfect duties. So, in the case of a perfect duty to return the borrowed money, the presence of options to pay by cash or by a bank transfer counts as latitude. Similarly, a choice among a range of courses of action that agents have when deciding how to comply with their duty of beneficence is also an instance of latitude permitted by morality.

⁵⁰ See Pummer (2016: 80–81) for a similar construal of cost.

⁵¹ Demandingness can also be understood in terms of difficulty rather than cost. I do not discuss it in this sense. See McElwee (2015) for an argument that demandingness as difficulty can ground a legitimate objection to a moral requirement and Cohen (2000: 171–174) for an argument against that view.

The latitude permitted by imperfect duties generally differs from that permitted by perfect duties in virtue of its extensiveness. So, imperfect duties tend to give more extensive latitude than perfect duties. We can understand the extensiveness of latitude in two ways: qualitative latitude and quantitative latitude.

The quality of extensiveness is determined by the number of dimensions along which latitude is permitted. For example, one such dimension is time. An agent could be permitted to choose the time at which to act to comply with the duty. If she has such a permission, she has some latitude along the dimension of *time*. If she is required to act at a specific time, she does not enjoy latitude along this dimension. There are other dimensions. For instance, if an agent is permitted to choose a type of action with which to act towards fulfilling her duty, she has latitude along the dimension of *act-types*. If her duty is directed or owed to others, she may have latitude regarding whom to act towards to, thus having latitude along the dimension of *recipients*. Finally, if an agent is permitted to choose among more or less costly acts, she has latitude along the dimension of *sacrifice*. There might be other dimensions.⁵² The point of this discussion, however, is not to offer an exhaustive list but to make clear the idea that latitude can be more or less extensive in the qualitative sense.

Latitude can also differ in its extensiveness in the quantitative sense. A duty permits more quantitatively extensive latitude than some other duty, if an agent may choose among more courses of action along a given dimension when fulfilling the former duty than when fulfilling the latter. For example, my duty to return the book that I have promised to return today at noon or tomorrow at nine in the morning (but not at any other times) plausibly permits less extensive latitude along the dimension of time than my duty of beneficence. After all, the former only permits two courses of action, whereas the duty of beneficence permits many more courses of action, provided that I will encounter many situations in which I will have an option of helping someone but where doing so will not be strictly required.

⁵² See discussions by Horgan & Timmons (2010: 34–35) and Rainbolt (2000: 243).

So, extensiveness of latitude can differ qualitatively or quantitatively. Imperfect duties tend to permit choosing among distinct courses of action along more dimensions (quality), and more courses of action along each of the dimensions (quantity), than perfect duties.⁵³ This is what I have in mind when I say that imperfect duties permit more extensive latitude than perfect duties. And the extensive latitude permitted by imperfect duties seems to be related to demandingness of morality in the sense of “limiting” it or “keeping” it at a moderate level. Latitude along the dimension of sacrifice seems most important for the claim about the relation between latitude and demandingness. It seems that if agents are permitted to choose among more costly and less costly courses of action, then morality is not overly demanding or at least not as demanding as it would be if such choice was not permitted.

My discussion in the remainder of this chapter aims to show that even if latitude permitted by imperfect duties does “limit” the demandingness of morality in some sense, this relation is not theoretically significant. As I shall argue, on any potentially significant interpretation of the relation, the claim about “limiting” is false and thus should be rejected, or question-begging, or trivially true and thus insignificant.

While I intend the result to cover all imperfect duties, my discussion focuses on the duty of beneficence as the primary example. Even if the result is limited to the duty

⁵³ My discussion does not imply that imperfect duties *always* permit more latitude than perfect duties. For example, the duty of gratitude may permit no latitude along the dimension of recipients since it is usually directed to specific persons. And if I make a promise to a group of people, it could be that I have latitude when choosing whom to act towards when fulfilling this promise. One may worry that this creates a difficulty in drawing the distinction between perfect and imperfect duties. But this worry is easily addressed. Difference in the extensiveness of latitude is not the ground of the distinction. Rather, it is a consequence of such a ground. For instance, on the Kantian view that I have introduced in Chapter 1, imperfect duties are distinguished from perfect ones on the basis of the fact that they prescribe ends to adopt and promote while perfect duties prescribe or proscribe acting in certain ways. The greater extent of latitude that imperfect duties permit would be a consequence of the fact that they prescribe ends. Since these ends do not specify precise courses of action that agents are required to take, they permit more latitude than perfect duties that do prescribe or proscribe courses of action.

of beneficence and does not generalise, it would still be significant since the question of demandingness is especially important in the context of our duty to help others.

3.3 The Two Interpretations of the Relation

Now, armed with the concepts of demandingness and latitude, we can formulate the two interpretations of the claim that latitude “keeps” morality from being too demanding or “limits” the demandingness. I shall first formulate these claims and explain why they could be theoretically interesting. After that, I shall argue that they fail to be theoretically interesting due to being false, question-begging, or trivial.

The relation can be interpreted as an *explanatory* one. On this interpretation, the latitude permitted by imperfect duties explains the reduced demandingness. So, when it is claimed that latitude permitted by imperfect duties “limits” the demandingness or “keeps” the morality from being too demanding, what is meant is that *because* imperfect duties permit extensive latitude, morality is not overly demanding. Note that the formulation I have just given is non-comparative. That is, the claim is that the presence of latitude explains why morality is not overly demanding rather than why morality is not as demanding *as it would have been without latitude*. But a comparative formulation should not be ruled out. It should not be ruled out because even if the non-comparative claim is false, the comparative need not be. It is a possibility that, for instance, latitude is the reason why morality is less demanding than it would have been without it, but that it is still very demanding in the absolute sense. And, as shall become clear, a separate argument is needed against the comparative claim. So it is worth considering both claims rather than just one. Thus, here is how the explanatory relation can be formulated as a pair of claims:

Explanatory Non-Comparative Relation (EXP-NC). If and because imperfect duties permit extensive latitude, morality is not overly demanding.

Explanatory Comparative Relation (EXP-C). If and because imperfect duties permit extensive latitude, morality is not as demanding as it would have been without this latitude.

If these claims were true, this would be a theoretically interesting result. Those attracted to moderately demanding views on morality or those who want to defend at least some limits on moral demandingness would only need to establish that there are imperfect duties that permit extensive latitude. This seems to give a new argumentative route for those arguing against the view that morality is extremely demanding.

A further reason to take the EXP-NC and EXP-C seriously is that some of the claims that are made in the literature on imperfect duties are ambiguous enough to allow for such an interpretation. While I would not attribute this view to any author specifically, since there is not sufficient evidence for that, there are claims that make the discussion of this view worthwhile. For example, Andrew Schroeder (2014: 25), while defending his account of imperfect duties, claims that “the motivation for pursuing an account of beneficence as an imperfect duty was to explain why individuals don’t bear the heaviest or most demanding burdens”. Consider also the claim by Greenspan (2010: 184) mentioned above: “...the notion of imperfect obligation is important to a common-sense deontological approach precisely as a limitation on morality’s binding demands”. This claim could also be interpreted as a statement of the explanatory connection (although I do not claim that this is the interpretation that Greenspan herself intended). Her own view on this matter is not clear from her paper.

The relation between latitude and demandingness can be construed in a weaker way, as a mere *implication*. Perhaps the extensive latitude does not explain the demandingness of morality, but merely implies it. So, while there are independent reasons for why morality is not too demanding, as long as imperfect duties permit extensive latitude, morality will also be not very demanding (or not as demanding as it would have been without latitude). It may seem questionable whether implication

without explanation is possible. As an example, consider an analogy from the debates about moral responsibility. It is often thought that the appropriateness of attitudes such as blame or praise as responses to certain behaviour on the part of some agent implies that this agent is responsible for the behaviour in question. However, many philosophers deny that the appropriateness of blame or praise explains why someone is morally responsible. The explanation plausibly rests on some further, independent facts about the agent, such as the facts that the agent knew the relevant details about the situation and that she had the relevant kind of control over her behaviour. But facts about the reactive attitudes are important: they may serve as evidence that an agent is or is not morally responsible for certain actions.

Something similar could be claimed about the extensive latitude permitted by imperfect duties if we accept that it merely implies certain limits on demandingness. Even if facts about the latitude permitted by imperfect duties do not explain why morality is not very demanding, they may serve as evidence for the limits on its demandingness. If we have good reasons for thinking that imperfect duties do permit such latitude, then we may draw conclusions about demandingness of morality on the basis of our beliefs about latitude. So, it is worth exploring whether latitude could merely imply certain limits on demandingness without explaining why these limits are in place. Here is how the relation can be formulated:

Conditional Non-Comparative Relation (CON-NC). If imperfect duties permit extensive latitude, morality is not overly demanding.

Conditional Comparative Relation (CON-C). If imperfect duties permit extensive latitude, morality is not as demanding as it would have been without this latitude.

A further reason to consider CON-NC and CON-C seriously is that even more claims in the literature can be interpreted as stating one of these claims. Greenspan's (2010: 204) claim that "...imperfect obligation ... can limit the moral demands on an

agent by leaving her leeway for choice” is relevant here, as is Rainbolt’s (2000: 249) claim that extensive latitude “keeps the obligation ... from being too binding, too demanding”. Furthermore, Daniel Statman (1996: 221–222) suggests that imperfect duties “somehow encourage a minimalistic approach to morality”. While he does not explicitly endorse a claim like CON-NC or CON-C, his discussion gestures in support of something close to these claims.

In the next two sections, I shall argue that none of the EXP-NC, EXP-C, CON-NC, and CON-C is of any theoretical interest. I shall first tackle the non-comparative claims, and then turn to their comparative counterparts.

3.4 Against CON-NC and EXP-NC

Let us consider CON-NC first. I argue that it is either false, or trivially true and thus theoretically insignificant. Here is the statement of this formulation once again:

Conditional Non-Comparative Relation (CON-NC). If imperfect duties permit extensive latitude, morality is not overly demanding.

3.4.1 When CON-NC is False

CON-NC does not cover all moral requirements, it extends only to imperfect duties. To have an interesting discussion about the truth of CON-NC we would have to assume that perfect duties are not extremely costly to satisfy anyway. If satisfying them could be extremely costly, then CON-NC could be false despite the extensive latitude permitted by imperfect duties. So let us, for the sake of the argument, assume that the costliness of complying with perfect duties is not an issue. This assumption should be held throughout the discussion of all the non-comparative and comparative theses.

If “imperfect duties permit extensive latitude” means that an agent may satisfy her imperfect duties by choosing among many alternative courses of action, then CON-NC is false. And this is so even if those alternative courses of action differ among one

another with respect to the cost of taking them. In other words, they differ with respect to latitude along the dimension of sacrifice. I shall present my argument using the following case as an illustration, call it *Twenty-Six Options*. Suppose that an agent has twenty-six alternative courses of action that differ from one another in exactly this way. Suppose also that taking any of these courses of action would be sufficient to satisfy the duty of beneficence. So, the agent then is permitted to choose from twenty-six alternative courses of action, where taking the first one has a cost a , the second $a+b$, the third $a+c$, and so on until $a+z$.⁵⁴ Each variable stands for a distinct amount of sacrifice with a representing the smallest amount and each consecutive variable after a representing a greater amount of sacrifice, with z representing the largest amount. The set of courses of action that the agent could take to comply with the duty of beneficence takes the form of a disjunction containing twenty-six disjuncts: “sacrifice a , or sacrifice $a+b$, ..., or sacrifice $a+z$ ”. Now, it is possible that taking any of these courses of action is extremely costly. It is not ruled out that a is such a substantial sacrifice that making this sacrifice is extremely costly. Thus, it is false that having extensive latitude even along the dimension of sacrifice implies that morality is moderately demanding. Note that the number of alternative courses of action can be increased from twenty-six to some larger number, but the conclusion would not change.

Two remarks are warranted to clarify this argument. It may appear that since one of the key aspects, in which the courses of action differ, is the amount of sacrifice involved, it may seem that I view the duty of beneficence as the duty primarily to make sacrifices. This need not be the case. The duty could primarily require promoting others’ well-being, and promoting others’ well-being merely *tends* to decrease one’s own well-being (since we spend resources on others rather than ourselves). Alternatively, the duty could primarily require adopting the end of others’ happiness (as it would on a broadly Kantian account). And the latitude it would permit would be the latitude in choosing different ways to cultivate the relevant attitudes and act on

⁵⁴ How the cost is measured is not important.

them accordingly. But it is also plausible that the twenty-six courses of action are just those different ways to adopt the relevant attitudes and promote the relevant end. There is no reason to think that these alternative courses of action could not differ from one another in their costliness.

The second remark is that imperfect duties are often thought to have indeterminate conditions of compliance, whereas in the example I have offered the conditions of compliance are precise: taking one of the twenty-six alternative courses of action is obligatory, while failing to take any of those courses of action is forbidden. It could be objected that when the conditions of compliance are indeterminate, agents have wide discretion: they are not required to choose courses of action that are determinately sufficient for complying with requirements of morality, but are also permitted to choose courses of action that are neither determinately sufficient nor determinately insufficient. And such courses of action might not be very costly thus rendering morality moderately demanding.⁵⁵

This is not a serious problem for my argument because if the example was reworked to have the duty with indeterminate conditions, it would work just as well. I have offered the example with precise conditions of compliance to make the key point of the argument as transparent as possible.

Let us consider the case of indeterminate conditions of compliance. As I have suggested in Chapter 2, a duty's conditions of compliance are indeterminate if and only if there are courses of action for which it is indeterminate whether taking them is sufficient for complying with the requirements. A Kantian imperfect duty of beneficence, for instance, has indeterminate conditions of compliance if it can be neither true nor false whether an agent does enough to count as having adopted the end of others' happiness. So, attempting to form some attitudes such as the intention to promote others' welfare and the attentiveness to others' needs might not be

⁵⁵ One could try to argue that while this does not render morality moderately demanding, it renders it less demanding than it would be without imperfect duties and their latitude. I consider these comparative claims further below.

determinately sufficient to count as having adopted the relevant end.⁵⁶ I have argued in Chapter 2 that regardless of whether a duty's conditions of compliance are indeterminate, there must be courses of actions taking which would determinately satisfy the requirements of this duty. And similarly, there must be courses of action taking which would be determinately insufficient for satisfying the requirements. If there are no such courses of action, the duty ceases to be action-guiding since it would not be possible to judge whether an agent complies with this duty or not. But it is a logical possibility that all moderately costly courses of action are such that it is determinately false that taking them would be sufficient for complying with the requirements of beneficence. And those courses of action that fall in the space of those that are neither determinately sufficient nor insufficient could all be very costly to take. Those that are determinately sufficient would then be even costlier. And if a course of action is determinately insufficient for complying with the duty, it is forbidden. Since all moderately costly course of action could be forbidden, morality would be very demanding even if the conditions of compliance were indeterminate (and they usually are in the case of imperfect duties).

Returning to *Twenty-Six Options*, it could be that even if the requirements of beneficence were indeterminate, all the twenty-six courses of action could be determinately insufficient for complying with the duty. So even if one were permitted to choose courses of action that would be neither determinately sufficient nor determinately insufficient for complying with the duty, all such courses of action could be very costly. That morality is moderately demanding does not follow from the fact that imperfect duties permit extensive latitude regardless of whether their conditions of compliance are determinate or not.

3.4.2 *When CON-NC is Trivially True*

My rejection of CON-NC has relied on interpreting “imperfect duties permit extensive latitude” as “one may satisfy imperfect duties by choosing from a wide range

⁵⁶ Because, for instance, one's attempts are not clearly wholehearted.

of courses of action”. The reason why CON-NC is false is that on this interpretation of “...extensive latitude” it is logically possible that all the courses of action that the duty permits taking are extremely costly to take. So, an interpretation of “...extensive latitude” that would render CON-NC true must be such that among all the permissible courses of action at least one is not extremely costly to take. On this interpretation, CON-NC can be reformulated as follows:

Conditional Non-Comparative Relation* (CON-NC*). If imperfect duties permit extensive latitude *such that they can be satisfied by taking a moderately costly course of action*, morality is not overly demanding.

On this formulation, CON-NC* is trivially true. Demandingness is understood in terms of costs. So, if moral requirements can be satisfied by taking moderately costly courses of action, morality is moderately demanding. It is then trivially true that if imperfect duties permit extensive latitude, morality is moderately demanding.

But this is not a significant result. It does not generate any theoretically interesting conclusions. That is because this new formulation of CON-NC* rests on building the limits on demandingness into the concept of extensive latitude. Clearly, if we claim that *what it means* to permit extensive latitude is to permit moderately costly options, we would always have the result that morality is not too demanding when we have extensive latitude thus understood. It would be the same as saying “from the fact that we may comply with moral requirements by taking moderately costly options it follows that morality is moderately demanding”.

A defender of the claim that CON-NC* is theoretically significant might challenge this line of reasoning by appealing to the parallel with the moral responsibility debate that I have drawn when introducing CON-NC. This defender could try to argue that even if CON-NC* is trivially true, it is still interesting. It is interesting in the same way the claim that facts about appropriateness of blame or praise imply facts about whether one is morally responsible for a certain action is. Even though blame or praise for an act are appropriate just in case and because the relevant agent is responsible for

this act, intuitions that this agent is apt to blame or praise could arise independently of facts that render this agent morally responsible. And so, these intuitions may serve as evidence that the agent is morally responsible even if we do not know the relevant facts (such as whether this agent had satisfied the relevant conditions of control and knowledge). Similarly, the defender of CON-NC* may argue that our intuitions that imperfect duties should permit latitude of certain extensiveness may serve as evidence that morality is not overly demanding even if we are unsure about how demanding it is.

I have two responses. First, it is not obvious that our intuitions about blame or praise are independent of the assumptions that we implicitly or explicitly make about the kind of control and the kind of epistemic state the relevant agent is in. So, it might be that we only have intuitions that someone could be appropriately praised or blamed because we hold beliefs about the facts that render this agent responsible. And then the intuitions do not *support* any view about whether the agent is morally responsible — they end up being a part of what is at question.

My second line of response is that even if intuitions about the appropriateness of praise and blame are independent from facts that render agents morally responsible, intuitions about latitude of imperfect duties are not independent from intuitions about demandingness. If we *define* permitting extensive latitude as permitting at least one moderately costly way of fulfilling imperfect duties, our intuitions about the extensiveness of latitude can't be anything else but intuitions about the level of demandingness. An intuition “the duty of beneficence permits extensive latitude” would be nothing but the intuition “the duty of beneficence can be satisfied by a moderately costly course of action”, but that is just a basic intuition about how demanding this duty is. And it is questionable whether such intuitions could be useful as evidence for the view that morality is moderately demanding when we are unsure how demanding it is. For in that case, we are not *supporting* any view on demandingness of morality, but rather stating an intuition that is part of what is being questioned. It should be concluded that CON-NC* is not a theoretically interesting or significant thesis.

So far, I have argued that CON-NC is either false, or trivially true (when it takes the form of CON-NC*). It is either false to say that extensive latitude “limits” demandingness of morality or it is trivially true and uninteresting. In the latter case, claiming that latitude “limits” demandingness would amount to a rather uninspiring claim that the presence of limits on demandingness (moderately costly options) limits demandingness.

3.4.3 *EXP-NC: False or Question-Begging*

Recall the stronger claim stating the non-comparative relation between latitude and demandingness:

Explanatory Non-Comparative Relation (EXP-NC). If and because imperfect duties permit extensive latitude, morality is not overly demanding.

It is stronger since, aside from the claim that extensive latitude implies that morality is not too demanding, it also includes the claim that the presence of such latitude explains the level of demandingness that morality has.

Now, if “imperfect duties permit extensive latitude” means “imperfect duties can be satisfied by a wide range of alternative courses of action”, EXP-NC is false. The reasons are the same as in the case of CON-NC. It is logically possible that imperfect duties permit a wide range of alternative courses of action, which differ in how costly they are to take, but that morality is still very demanding because each of those courses of action is very costly to take. And if a weaker thesis, CON-NC, is false, then a stronger thesis, EXP-NC, is false too. What could explain moderate demandingness is not the presence of a wide range of alternative courses of action, but the presence of courses of action that would be moderately costly to take.

Suppose then that being permitted to take moderately costly courses of action is a necessary part of having extensive latitude, just as we have assumed to formulate CON-NC*. We can then formulate the analogous version of EXP-NC:

Explanatory Non-Comparative Relation* (EXP-NC*). If and because imperfect duties permit extensive latitude *such that they can be satisfied by taking a moderately costly course of action*, morality is not overly demanding.

Unlike CON-NC*, EXP-NC* is not trivially true. However, it is question-begging. Recall that CON-NC* amounts to the claim that if imperfect duties can be satisfied by taking a moderately costly course of action, then morality is not overly demanding. It would be question-begging to claim that the fact that the former claim explains the latter because the former claim presupposes the latter. The claim “imperfect duties permit latitude that is extensive in the sense that they can be satisfied by a moderately costly course of action” presupposes the view that morality is not overly demanding rather than explains it. To claim that imperfect duties permit extensive latitude thus understood, we need to rely on a substantive view about demandingness of morality or the availability of moderately costly options for satisfying moral requirements.

The argument I have just made can be illustrated with an example. On a broadly Kantian account of imperfect duties, latitude is the consequence of the fact that obligatory ends do not prescribe a uniquely specified course of action. Many distinct courses of action are consistent with successfully adopting and promoting the obligatory ends. But to claim that at least one of these courses of action must be not costly, we would need to make and defend substantive assumptions about demandingness of the relevant duties, and plausibly of morality as a whole.

So, if extensive latitude is interpreted as it is in EXP-NC*, then instead of explaining moderate demandingness, extensive latitude presupposes it.

In this section, I have argued that neither CON-NC nor EXP-NC is of any theoretical significance. On a standard interpretation of “extensive latitude” as permissibility to satisfy imperfect duties by choosing from a wide range of courses of action, they are both false. If the interpretation of “extensive latitude” is modified as to guarantee that at least one moderately costly course of action is among the relevant courses of action, then the appropriately modified versions of CON-NC and EXP-NC, namely CON-NC* and EXP-NC* are trivially true and question-begging respectively.

Let me say a bit more about the distinction between trivial truth and question-beggingness in this context. EXP-NC* is begging the question because it presupposes what needs to be explained in the explanation. EXP-NC* purports to explain why morality is moderately demanding. But all it gives as an answer is that morality is moderately demanding because its requirements can be satisfied by taking moderately costly actions. This begs the question: “Why can they be satisfied by taking such actions?” By contrast, CON-NC* does no such thing. CON-NC* is the claim that if the requirements of morality can be satisfied by moderately costly actions, morality is moderately demanding. This is an uninteresting trivial truth. But it is not question-begging because it is not intended as an explanatory or a justificatory claim.

3.5 Against CON-C and EXP-C

Let us now consider the comparative theses about the relation between latitude and demandingness, CON-C and EXP-C. The argument follows roughly the same path as the argument against CON-NC and EXP-NC. However, some specifics are different. Moreover, as I have mentioned above, even though the non-comparative claims are false, the comparative ones could still be true. So, it is worth discussing them explicitly. I shall argue that CON-C is either false or trivially true, while EXP-C is either false or question-begging.

Here is a statement of CON-C once again:

Conditional Comparative Relation (CON-C). If imperfect duties permit extensive latitude, morality is not as demanding as it would have been without this latitude.

3.5.1 *When CON-C is False*

Suppose that imperfect duties can be satisfied by taking one of many alternative courses of action. Assume, moreover, that these alternative courses of action differ

from one another with respect to cost of taking them (among other possible differences). Now, would a duty that permitted less extensive latitude than the one just described be more costly to comply with? And would it follow that morality is less demanding if it included a duty that permitted more extensive latitude rather than a duty that permitted less extensive latitude? I argue that both questions should be answered negatively. The answers depend on how the counterfactual should be interpreted. Specifically, it is important how duties that “permit less extensive latitude” should be understood. Recall that latitude can differ in its extent qualitatively and quantitatively. Suppose that we do not restrict the space of acceptable counterfactuals. So, any duty that permits less latitude *in any way at all* — that is, qualitatively or quantitatively — could be part of the counterfactual. In such a case, CON-C would be false.

Consider *Twenty-Six Options* again as an illustration. Now, imagine that the duty can be satisfied not by any of the twenty-six courses of action but only by any of the twenty-five, because $a+z$ is no longer a permissible option.⁵⁷ The fact that one duty permits more extensive latitude does not imply that this duty is less costly to satisfy. This is because the least costly course of action is the same in both cases, it is a . Thus, it does not follow that morality would be less demanding if imperfect duties permitted less extensive latitude than they could otherwise permit. So, CON-C is false. And this result would hold regardless of the specific kind or extent of latitude permitted, as long as there are no restrictions on what can count as the relevant counterfactual.⁵⁸

⁵⁷ For instance, if $a+z$ involves making a sacrifice on the part of the agent so great that it negatively affects the relationship between the agent and the recipient of the benefit and makes this relationship grossly unequal (Herman 2021: Ch. 3).

⁵⁸ As in the case of CON-NC, making the conditions of compliance indeterminate would not render CON-C true.

3.5.2 When CON-C is Trivially True

There is a way of imposing limits on what could count as the relevant counterfactual in a way that would render CON-C true. Duties that permit less extensive latitude should be understood not as duties that permit less extensive latitude in any way at all. Rather, they should be understood as duties that permit latitude such that the least costly course of action sufficient to satisfy them is costlier than the least costly action sufficient to satisfy the duty with more extensive latitude. Let's consider the above example from *Twenty-Six Options* again. Now, the duty that permits twenty-five courses of action would not be part of the relevant counterfactual because the least costly way of satisfying it is by taking *a*. And *a* is also the least costly way to satisfy the duty with more latitude, the duty with twenty-six courses of action. We have seen above that in this case, CON-C is false precisely because *a* is the least costly option sufficient for satisfying both duties.

But now let us consider a duty that could be part of the newly restricted relevant counterfactual. Suppose that we compare the duty with more latitude to a different duty that permits twenty-five courses of action. But now this second duty cannot be satisfied by taking *a*. It can be satisfied by *a+b* and any further course of action up to *a+z*. Now, this duty is clearly costlier to satisfy than the one permitting twenty-six options, since the least costly way of satisfying it, *a+b*, is costlier than *a*, which is sufficient to satisfy the duty with more extensive latitude. If what can count as an acceptable counterfactual were restricted in this way, then CON-C would be true. However, trivially so.

By restricting the space of acceptable counterfactuals this way, we would be effectively claiming that “duties that permit less extensive latitude” is equivalent “duties that are costlier to satisfy”. Conversely, we would be claiming that “duties that permit more extensive latitude” just is “duties that are less costly to satisfy”. CON-C could then be reformulated:

Conditional Comparative Relation* (CON-C*). If imperfect duties permit extensive latitude *such that they can be satisfied by a less costly option than duties*

permitting less extensive latitude, then morality is not as demanding as it would have been without this latitude.

It should now be clear that CON-C* is trivially true. It is true because it presupposes that “having less costly options to satisfy the duty” is part of the meaning of “having more latitude to satisfy the duty”. But, similarly to the case of CON-NC*, this is not an independent intuition that could be used as evidence that morality is not as demanding as it would be without extensive latitude. This is merely a restatement of the intuition that morality would not be as demanding without extensive latitude. It does not in any way support this intuition. CON-C* is not a significant or interesting thesis.

3.5.3 EXP-C: False or Question-Begging

Here is a reminder of how EXP-C is stated:

Explanatory Comparative Relation (EXP-C). If and because imperfect duties permit extensive latitude, morality is not as demanding as it would have been without this latitude.

I take it that little needs to be said at this point to show that EXP-C is false, if by saying that imperfect duties permit extensive latitude it is meant that there are many alternative courses of action such that taking any one of them is sufficient to satisfy them. The availability of many alternatives would not explain the lower demandingness. Only the availability of less costly alternatives would.

But if we assume that less costly alternatives are permitted and claim that this explains why morality is not as demanding as it would be without such alternatives, we run into the problem of begging the question. We run into this problem because to support the claim that less costly alternatives are permitted, we need to make and defend assumptions about demandingness of morality. We, therefore, cannot use the

claim about the permissibility of less costly alternatives to explain claims about demandingness.

3.6 How Should the Relation Between Latitude and Demandingness be Construed?

I have now argued against all the four claims, CON-NC, CON-C, EXP-NC, and EXP-C. The former two claims are either false or trivially true. The latter two are either false or question-begging. If the claim that extensive latitude permitted by imperfect duties “limits” demandingness or “keeps” morality from being overly demanding is understood in any of these four ways, then it is at best misleading and at worst false. But we should not think that the philosophers who made claims such as these are all guilty of some serious mistakes. There is a way to interpret their views more charitably. But even the most charitable interpretation of the relation does not render it very significant or interesting.

Here is how, I think, the relation between latitude and demandingness should be understood:

Non-Comparative Compatibility Relation (NC-CR). If imperfect duties permit extensive latitude, it is logically possible for morality to be moderately demanding.

Comparative Compatibility Relation (C-CR). If imperfect duties permit extensive latitude, it is logically possible for morality to be not as demanding as it would have been without this latitude.⁵⁹

⁵⁹ There are no restrictions on what may count as an imperfect duty without permitting extensive latitude. After all, this thesis is quite weak. The reason why morality is not as demanding could be independent of why it includes imperfect duties with extensive latitude.

This relation is even weaker than the previous two. Even if extensive latitude neither explains nor merely implies the limits on demandingness, it could at least be compatible with them. So, including imperfect duties with extensive latitude in our moral theorising does not rule out the logical possibility of morality being not too demanding (or being not as demanding as it would have been without the imperfect duties and their latitude).

Given the weakness of this relation, it is difficult to deny that it holds. Its truth is also not trivial: we need not presuppose any view on demandingness to claim that the permissibility of choosing among many courses of action is compatible with morality being not too demanding. A similar claim can be made about the comparative relation.

The compatibility relation also offers a plausible picture of what hides behind the claims such as the one that latitude “limits” demandingness. We start with the motivation to give an account of morality on which it would not be extremely demanding. Introducing the distinction between perfect and imperfect duties such that the latter permit significantly more extensive latitude creates a nice framework, in which the moderately demanding view on morality can be defended.⁶⁰ But by themselves, facts about latitude permitted by imperfect duties have no interesting implications for the facts about demandingness of morality.

3.7 ‘Confinement’ as Demandingness?

I shall now consider a potential objection to my argument. So far, I have proceeded on the assumption that demandingness should be understood in terms of costs of satisfying moral requirements. I have also implicitly assumed that the cost of complying with requirements of a duty is just the cost that an agent incurs by taking a course of action to satisfy the requirements. Some, however, may disagree and suggest that there is an additional dimension to the cost of satisfying moral requirements which is independent of cost of taking any specific course of action, namely, ‘moral

⁶⁰ Noggle (2009) can be argued to be engaging in a project of this kind.

confinement'.⁶¹ Moral confinement can be understood as the opposite of latitude. The more courses of action or options such that any one of them can be taken to satisfy the requirements of a duty are available to an agent, the more latitude this agent enjoys. The fewer options of this sort are available to an agent, the more she is morally confined. The idea is that being more morally confined is more costly for an agent than to be less morally confined, other things being equal. So, if a duty gives an agent more options for satisfying its requirements than it would in some counterfactual circumstances, then it is less morally confining and, other things being equal, less demanding just in virtue of permitting more options, regardless of how costly it is to take any of these options.

The idea that moral confinement is a species of demandingness may pose problems for my argument by giving direct support to claims such as CON-C. Recall that CON-C is the claim that if a duty permits extensive latitude, then it is not as demanding as it would be without permitting such latitude. I have argued that it is false or trivially true. But if moral confinement is a kind of demandingness, then there is a new reason to accept claims like CON-C. The more extensive latitude a duty permits, the less confining, and thus the less demanding it is.

Why would one think that moral confinement is costly? Note that this cost cannot be reduced to the costs of taking any given option. Rather, the cost must be grounded in the lack of options regardless of the cost of taking them. Liam Murphy (2000: 30–31) puts the key idea behind the objection from moral confinement as follows: “An extremely confining moral theory conflicts with a certain ideal of the moral agent. It is simply morally unappealing, it could be said, to think that moral agents have their lives fully mapped out for them by the dictates of morality. Much more appealing is a picture of moral agents as having broad space for individual decision within limited

⁶¹ Suggestions along these lines are present in van Ackeren (2018), Scheffler (1992: 98), and in Shiffrin (1991: 250–251). Liam Murphy (2000: 26–28) also discusses whether confinement can be argued to make a moral theory more demanding. I borrow the term “confinement” from Murphy (2000: 26). To confine an agent morally would be to infringe upon her “moral autonomy₂” in Shiffrin’s terms.

constraints”.⁶² So, when morally confined, agents are deprived of an opportunity to make choices among morally permissible alternatives. And since this opportunity is valuable, being deprived of it is costly.

Why is this opportunity valuable? One reason that is briefly mentioned in Murphy’s characterisation of the objection and sometimes invoked in the literature, is that it is valuable to be the agent who defines and pursues her own conception of the good life, and doing so is impossible for an agent who has only one morally permissible course of action that she is required to follow throughout her life. Seana Shiffrin (1991: 251) proposes something similar: “Within fully directive theories, agents are deprived of the substantive opportunity to create a unique identity and to engage and channel their creative and deliberative powers and energies upon themselves and their interests, just for their own sake”.⁶³

I think that this reasoning is mistaken. Consider the following pair of cases to see why.

Good Life – Options (GLO). Morality imposes some requirements on John – such as to promote others’ welfare to an extent x , respect others’ rights, and improve himself in some moral and non-moral respects – but permits significant latitude of different kinds. For instance, John may choose whom to help, when to help, in what way to help, provided that he promotes the good to the extent x . He also may choose which capacities and talents to develop. There are many acts that John may choose that are neither required nor prohibited but merely permissible. Given that morality permits an extensive range of options, John has an opportunity to find the options that correspond to his view of the good life and pursues

⁶² Although Murphy, as most other people in the literature, discusses confinement in the context of moral theories rather than individual duties, I think that the point he makes is also relevant if we consider confinement just in the context of the duty of beneficence.

⁶³ See also Brock (1991: 910–911) who puts forward a very similar idea.

such a life by taking these options. He complies with the requirements of morality while also enjoying valuable relationships with others, engaging in creative work, and successfully carrying out several other important personal projects.

Good Life – Confinement (GLC). Morality requires John to promote others' welfare to an extent x , respect others' rights, and improve himself in some moral and non-moral respects but permits virtually no latitude. John is morally required to help specific people, on specific occasions, in specific ways, and is required to develop certain capacities and talents. There are no merely permissible acts, and every act is either required or prohibited. However, despite the absence of options, the course of actions that morality prescribes perfectly coincides with the course of action that John himself would choose while pursuing his conception of the good life. John complies with the requirements of morality while also enjoying valuable relationships with others, engaging in creative work, and successfully carrying out several other important personal projects.

So, John pursues the same course of action in both cases, and this course of action corresponds to his view of the good life. However, in the latter case he also has decisive moral reasons to pursue this particular course of action, while he lacks such reasons in the former case. The point for designing these two cases is to hold fixed the levels of demandingness that is not due to confinement, so we can better test the intuitive plausibility of the idea that confinement is a species of costs *distinct from the ordinary ones*. Since, by hypothesis, these two cases are identical in terms of non-confinement demandingness, if the confinement view is plausible, then we should judge these cases to be morally asymmetric. But it is not clear that we judge them like that. And even if we do, this is better explained not by the importance of confinement, but by other factors.

So, the cases are supposed to show that it is at least possible for an agent to pursue her conception of the good life while being morally confined to a single course of action. That is possible if one is morally required to lead a life that one would lead for non-moral reasons even if it was not morally required. Given these two cases, it is less clear to me whether there is any additional cost that an agent incurs solely in virtue of being morally confined. If the cost is due to one's lack of opportunity to pursue one's conception of the good life, then it seems that the problem is not in the lack of choice among morally permissible alternatives but in the lack of specific courses of action, which may be permissible and available to the agent even if there are no other morally permissible options. Furthermore, it is unclear to me whether the lack of permissible alternatives is costly for John if the course of action that he is morally required to take is also the course of action that makes things go best for him from the perspective of his well-being. Does the introduction of alternatives, taking which would be worse for John, non-trivially reduce the costs he incurs?

The defender of the objection from moral confinement may respond that what is problematic is that John does not really choose to pursue his conception of the good, and that choice, that exercise of one's autonomy, is valuable.⁶⁴ If morality does not give much latitude to agents, it is costly precisely because it deprives them of the value of making the choice. An analogy with paternalism may seem apt. Even if one is paternalistically coerced to do what is in one's best interests, there is something problematic about paternalism. Similarly, even though John is morally required to do what he would choose for non-moral reasons anyway, there is something problematic about his not exercising autonomous choice.

My response to this is that John is not deprived of an opportunity to exercise his autonomy in GLC. He is free to judge and weigh reasons for and against the alternative courses of action available to him. He may recognise that one of the courses of action is supported by a decisive moral reason, but he is still free to disregard that reason and

⁶⁴ For a related point, see Shiffrin (1991: 252).

act against it.⁶⁵ Note that in GLC he has no self-interested reasons for acting against the moral reason, so doing that would also be imprudent. I struggle to see the value of having more than one morally permissible option in cases where one (1) has an opportunity to exercise one's autonomy and decide how to act and (2) where the course of action supported by a decisive moral reason is the one that makes things go best for one from the perspective of one's well-being.

Furthermore, the analogy with paternalism is misguided. An immediate reason why this is so is that treating people paternalistically implies being able to treat them in any way. And morality just does not have the requisite agential capabilities. But even if the claim that confining morality is paternalistic is understood metaphorically, the analogy still would not work. Paternalism is problematic not because it is costly for an agent to be deprived of choice. The point of paternalism is precisely to reduce costs on people by making them do what is in their interests. The problem with paternalism, very briefly, is that it does not respect people's rights. So, if one claims that a moral theory is implausible because of the costs it imposes on agents, this implausibility is not explained by the reasons that explain why paternalism is morally objectionable. The analogy with paternalism suggests that rather than being problematic because of the costs it imposes, a moral theory is problematic because it does not recognise agents' rights or entitlements. But saying that a moral theory does not recognise certain entitlements implies a rejection of this theory and a commitment to a different moral view which affirms those entitlements. So, to criticise a moral theory for this would be to say that this theory is mistaken due to its failing to recognise the relevant entitlements regardless of the costs that it imposes on agents.⁶⁶

⁶⁵ Kagan (1991: 922–923) makes a similar point.

⁶⁶ Murphy suggests something similar (2000: 30). Theron Pummer (2023: 27–30) takes a similar route when defending autonomy-based permissions. Pummer's idea seems to be that there is a reason to reject highly confining versions of morality even if these versions were not too demanding. The reason to reject these versions is that they do not recognise the autonomy of agents. This is consistent with my response to the objection in this section. All that I claim is that moral confinement is not a species of cost, and so a highly confining version of morality might not be more demanding than a non-confining

I have been unable to find a good account of why moral confinement is costly if the costs are independent from costs of lacking specific options. But even if it is granted that John is somewhat better off in GLO than in GLC because in the former he has more morally permissible options, it is unreasonable that the cost of lacking such options could be appealed to in order to claim that what morality requires in GLC is too demanding. Although morality in GLC may seem unappealing this seems to be not because of the costs it imposes on John, but for other reasons. It could be, for instance, that moral facts do not always yield a division of all acts into required and prohibited, and such a complete division may only result from mistakenly thinking that some morally irrelevant features are morally relevant. But then the problem with morality in GLC is not the costs it imposes on agents. Thus, even if we judge these cases to be morally asymmetric, the asymmetry is not due to differences in demandingness.⁶⁷

Since moral confinement itself is not costly, or at least, not costly in a way that matters for the question of demandingness, the objection from moral confinement does not support claims such as CON-C and does not threaten my argument.

one (and even if it would be more demanding, it would not be because of its being more confining). But it may very well be that there are non-cost-based reasons to accept a non-confining version of morality.

⁶⁷ One might suggest a further counterargument. If morality requires each person to follow a uniquely specified course of action, figuring out what exactly one is required to do might be costly in terms of cognitive effort. It might be costly to always try to determine precisely what one is required to do. Morality that permits latitude might be less demanding, the objection would go, since it would not require as much cognitive effort on the part of agents. I do not find this point sufficiently persuasive. While one would have to try to figure out exactly what is required if morality was confining, one would also have clear guidance as to what one should do. If morality was not confining, choosing among a large number of permissible options might be no less costly than trying to figure out what exactly is required.

3.8 Conclusion

My aim in this chapter has been to argue that there is no theoretically interesting relation between extensive latitude permitted by imperfect duties and the demandingness of morality. I have argued that the relation should be understood as one of mere compatibility. This relation is not very interesting: morality could be moderately demanding while including imperfect duties, but it needn't be so. Nevertheless, this is a welcome result: it helps us make sense of what different theorists must mean when they claim that latitude "limits" the demands of morality. They must mean that imperfect duties are part of a general framework, in which a moderately demanding version of morality can be defended. But the fact that imperfect duties and the latitude they permit are part of this framework does not lend any support to the claims about demandingness. It is valuable to realise that questions such as "how much latitude does the duty of beneficence permit?" and "how demanding is the duty of beneficence?" are importantly different and answering one of them may have no implications for answering the other.

Chapter 4

Imperfect Duties and Collective Impact Cases

4.1 Introduction

This chapter discusses the inefficacy problem that arises in collective impact cases and a specific solution to this problem. Collective impact cases are such that when enough agents act in certain ways, a harmful or beneficial outcome occurs, but no individual agent's actions seem to make a difference to this outcome.⁶⁸ Consider a simple example:

Fishermen. People in a town survive by fishing in a nearby lake. If all fishermen use small nets to fish, everyone catches just enough fish to feed themselves and their families. If some sufficiently small number of fishermen use big nets, these fishermen reap significant benefits of catching more fish than everyone else, while most fishermen who use small nets still catch just enough fish to survive. If too many fishermen use big nets, then the fish population in the lake declines substantially and as a result the people in the town starve.⁶⁹

⁶⁸ Derek Parfit (1984) is probably responsible for the popularity of debates about such cases. Another classic discussion is by Shelly Kagan (2011).

⁶⁹ Parfit (2011: 303–304) discusses a similar case.

Collectively, many fishermen using big nets cause starvation. However, it is not clear whether any individual fisherman is making a morally significant contribution to this outcome by choosing to use a big net to fish on any given occasion. Intuitively, it seems that at least some of the fishermen should not use big nets, that it is morally wrong for them to use big nets. However, it is tricky to explain why any of the fishermen ought not to do so, given that a single fisherman's using a big net on any given occasion does not seem to make things go any worse. The inefficacy problem is exactly this – the problem of identifying what it is that an individual agent is doing wrong when acting in collectively harmful ways and explaining why it is wrong given that her so acting has no morally significant impact. By “collectively harmful ways”, following Julia Nefsky (2021: 211), I understand ways such that if enough people so act, the harmful outcome occurs.⁷⁰

There are three kinds of response to the inefficacy problem (Nefsky 2019). First, there are sceptical responses to the problem.⁷¹ These responses claim that it is difficult to explain why individual agents acting in collectively harmful ways are doing anything wrong because they are not in fact doing anything wrong. Even though there may be something attractive about these responses, we should not assume that they are correct. It is worthwhile to look for a non-sceptical solution.

Second, there are solutions that deny that individual agents who act in collectively harmful ways do not make any morally significant difference.⁷² These solutions require us to accept that even though it may seem that individual acts do not seem to make a morally significant difference, at least one act must make such a difference. These cases are often called ‘threshold cases’ because in these cases, a sufficient number of acts of the relevant type trigger a morally significant effect upon hitting a precise threshold. For the purposes of this chapter, I shall not assume that all

⁷⁰ While my discussion should extend both to cases of collective harm and collective benefit, I shall stick to talking about the former, if only for the sake of brevity. I shall refer to collective impact cases with a potentially harmful outcome as “collective harm cases” or “cases of collective harm”.

⁷¹ For instance, Kingston & Sinnott-Armstrong (2018).

⁷² Arntzenius & McCarthy (1997), Hiller (2011), Kagan (2011), Otsuka (1991), Regan (2000).

collective impact cases are threshold cases. On the contrary, I shall assume that there are cases such that for any act of the relevant kind, there is no chance that this act by itself triggers a morally significant effect. Cases of this kind are called ‘non-threshold cases’. So, it is possible that no instance of using a big net in *Fishermen* makes a morally significant difference to the outcome. In other words, it is possible that *Fishermen* is a non-threshold case. This is not an implausible assumption, given that there is no decisive reason to believe that non-threshold cases are impossible (Nefsky 2011).⁷³

Third, there are solutions which deny that making a difference to an outcome is necessary for moral wrongness. On this view, even if none of the individual acts of the relevant type makes any morally significant causal difference to the outcome, this does not entail that individual agents performing these acts are not doing anything wrong.⁷⁴

In this chapter, I focus on a specific kind of solution from the third group. This solution claims that acting in collectively harmful ways is wrong because it violates a certain kind of imperfect duty. Specifically, I consider the end-based Kantian account of imperfect duties that I have outlined in Chapter 1. The solution that I shall examine holds, roughly, that when acting in collectively harmful ways, individual agents fail to take others’ happiness or welfare as an end, and that amounts to violating the corresponding imperfect duty. My aim is to argue that in most cases, acting in collectively harmful ways would not violate any imperfect duties, and thus the proposed solution fails. The main reason for this is that when trying to address the inefficacy problem, potential solutions restrict the scope of assessment to agents’ patterns of activity in one isolated case (e.g., *Fishermen*). But assessing such patterns in isolation makes it impossible to make any judgments regarding the agents’ compliance with imperfect duties. However, if the scope of assessment is widened to

⁷³ In the next chapter, I shall argue that there are reasons to believe that non-threshold cases at least of a purely phenomenal variety are impossible.

⁷⁴ Albartzart (2019), Parfit (2011).

make such judgments possible, it is overwhelmingly unlikely that agents violate any imperfect duties.

My discussion proceeds as follows. In section 4.2, drawing on Nefsky's (2015, 2021) arguments, I formulate two desiderata for a solution to the inefficacy problem, and in section 4.3 outline the Kantian account of imperfect duties, which seems to fit the desiderata well. In section 4.4, I offer some general reasons for thinking that Kantian imperfect duties do not help in solving the inefficacy problem. In section 4.5, I outline Maike Albertzart's (2019) specific attempt to address the inefficacy problem by an appeal to this view on imperfect duties. In section 4.6, I argue against this attempt.

4.2 The Desiderata for a Solution

In this section, I present the two main desiderata for a solution to the inefficacy problem. I would like to make it explicit that the plausibility of my main argument in this chapter does not depend on whether we accept these desiderata. However, its significance does. The imperfect-duty-based solution that I discuss in sections 4.3 and 4.4 seems attractive, at least in part, because it satisfies these desiderata. And it thus would be more significant to show that this solution is not available. But if the desiderata are rejected, then there is less reason to consider this solution in the first place, which would make the argument against this solution somewhat less exciting. I believe, however, that the desiderata are plausible.

4.2.1 *Identifying the Desiderata*

One desideratum for a solution is that it identifies what it is that individual agents in collective harm cases are doing wrong.⁷⁵ In cases like *Fishermen*, we want to be able to justifiably claim that at least some individual fishermen are doing something wrong

⁷⁵ It is of course also important for the solution to explain why the identified activity is wrong. I shall not mention this further on.

when they act in ways that collectively cause starvation. One straightforward way to understand this desideratum is to construe it as the claim that an acceptable solution should identify facts that make acting in collectively harmful ways *pro tanto* wrong, that is, wrong without a sufficiently strong justification for so acting. There is a good reason to think that at least in some cases this is what we want from a solution. Consider the case of voting in a one-off election in a simple majority system. In such a case, each person has only one vote. Suppose that if not enough votes are cast for the candidate A, the candidate B is elected, and this leads to a significantly harmful outcome. Now, if failing to vote for A is not *pro tanto* wrong, then it seems that the solution does not show that agents are doing something wrong when they collectively cause harmful outcomes. If the only thing that a voter could possibly do with respect to the harmful outcome in this case is vote and if failing to vote is not even *pro tanto* wrong, then if people fail to elect A, it seems that we cannot find fault in anything that they have done.⁷⁶ We would then be forced to conclude that collectively causing a harmful outcome does not involve any wrong actions on any individual agent's part. To avoid conclusions of this kind, we should accept the desideratum that *an acceptable solution must be able to classify individual acts of the relevant kinds as pro tanto wrong in some cases.*

The first desideratum is not all that we want from a solution. There is good reason to think that while we want individual acts of the relevant kinds to be classified as *pro tanto* wrong in *some* cases, we want to avoid this result in others. Nefsky (2021) argues that solutions that solve the problem by identifying a fact that makes *any* single act of the relevant kind *pro tanto* wrong would be implausibly strong. Such solutions would classify intuitively permissible acts as *pro tanto* wrong. Suppose that anthropogenic climate change is a collective harm case, where acts that are collectively harmful are various CO₂ emitting activities. One such activity is driving. It is plausible to assume

⁷⁶ Similar considerations apply to Parfit's (1984) cases such as the *Harmless Torturers*. If pushing the button is not even *pro tanto* wrong, we seem to be forced to conclude that none of the torturers is doing anything wrong when they collectively cause the victims to experience pain.

that one person's going for a short drive has a very little chance of making things worse with respect to climate change-related harms, while if many people drive regularly, they collectively make a considerable difference for the worse.⁷⁷ Intuitively, we must be able to find some fault in the drivers' activities. However, claiming that driving is *pro tanto* wrong (that is, wrong without a strong justification) seems implausibly strong. It seems that, other things being equal, an agent does not need a moral justification for going on a short drive on some single occasion, given how unlikely it is to make a moral difference. Going for a single drive just because it is convenient or pleasant is intuitively not *pro tanto* wrong. We should then accept a second desideratum, namely, that *the solution should not classify collectively harmful acts as pro tanto wrong in all cases*. In some cases, such as the climate change one, while drivers are doing something wrong when they collectively make a significant contribution to climate change-related harms, it is not because each individual instance of driving a car is *pro tanto* wrong. Similarly, if we return to *Fishermen*, it is counterintuitive that each instance of using a big net is *pro tanto* wrong.⁷⁸ So, an acceptable solution should find other ways to condemn what fishermen are doing.

Nefsky's (2021: 212–215) own proposal is that to see whether an individual agent is doing anything wrong, we should look at her overall pattern of acting in collectively harmful ways, rather than single actions. One unnecessary drive or using a big net once are not *pro tanto* wrong but engaging in these activities often and without good reasons seems clearly wrong. The upshot of Nefsky's proposal is that in many cases, it is not *pro tanto* wrong to act in collectively harmful ways, but it is wrong to act in these cases too much of the time.⁷⁹ Why acting in collectively harmful ways too much of the time

⁷⁷ Let's abstract from other significant sources of CO₂ emissions such as burning of fossil fuel for electricity and heat in this case.

⁷⁸ If the example seems counterintuitive for reasons related to animal welfare, it can be modified such that no animals are harmed. Substitute the fishermen for people who rely on the lake as their source of water, where using big buckets too much of the time would lead to depletion and drought.

⁷⁹ In Nefsky's (2021) terminology, having this feature already makes an obligation imperfect. That is, an obligation is imperfect insofar as it requires acting in certain ways enough of the time or avoiding

is wrong and what counts as “too much” would vary depending on the specifics of the solution.⁸⁰

Thus far, we have identified the two desiderata:

(1) Acts of the relevant kinds must be classified as pro tanto wrong in *some* cases.

(2) Acts of the relevant kinds must not be classified as pro tanto wrong in *all* cases.

The second desideratum entails that in some cases, there must be other ways of condemning what the agents do. For instance, classifying their patterns of activity as wrong (as opposed to specific kinds of acts).

4.2.2 *Clarifying the Desiderata*

How do we distinguish cases in which we should classify individual actions as pro tanto wrong from those in which we should not? The following considerations that are implicit in the above discussion may serve as a rough guide. In cases where failing to classify specific acts of the relevant kinds as pro tanto wrong would imply failing to find anything wrong about the acts that lead to collectively harmful outcomes, these specific acts should be classified as pro tanto wrong. In cases where failing to classify specific acts as pro tanto wrong would not imply failing to find something wrong about the acts that lead to collectively harmful outcomes, these specific acts should not be classified as pro tanto wrong. So, in one-off cases like voting, a single act of the relevant kind should be plausibly classified as pro tanto wrong. By contrast, going for a single drive or using a big net once should not be classified as pro tanto wrong (instead, patterns of activities involving these acts could be classified as wrong). For this distinction to be defensible, there must, of course, be a deeper reason for distinguishing these kinds of cases. I shall not discuss what kind of reason that might

acting in certain ways enough of the time. To avoid confusion, I do not adopt this terminology. The way I use the term is consistent with many obligations regarding one’s choices over time being classified as perfect.

⁸⁰ Driving too often might be wrong, for example, because it has a higher chance of contributing to climate change-related harms, or it might be because by driving often the driver exceeds her fair share of emissions. It might be some combination of the two (Baatz & Voget-Kleschin 2019).

be. For my purposes, it is sufficient that the rough guide I have just described helps sort the cases in an intuitive way.

There is another question in the background. Why do we need to find wrongness in the actions of agents in collective harm cases? Without answering this question, it would be difficult to find the two proposed desiderata plausible. The answer is that it seems that if agents collectively cause a harmful outcome, while avoiding this outcome was a reasonable option, it is plausible that at least some of the agents did something wrong. This idea is sometimes referred to as the “principle of moral harmony”.⁸¹ The principle can be put as follows. If we all always do what morality requires of us, then we would always bring about the morally best outcomes which we have the option of collectively bringing about.⁸² In *Fishermen*, for instance, it seems clear that the fishermen have the option of keeping the town fed while avoiding famine. And this would be a better outcome than causing the town to starve. So, it is plausible that if the fishermen bring about this worse outcome that they could avoid, then at least some of them have done something wrong.

I have one final remark about the two desiderata. I should emphasise that to satisfy the desiderata, it is necessary to identify features that make it wrong to *act* in collectively harmful ways. Condemning the agents on the grounds that have nothing to do with their actions that together cause harm would amount not to solving the problem, but to avoiding it. For example, arguing that fishermen have a duty to come up with an institutional arrangement that would prevent overfishing would not solve the problem. It would not solve the problem because it is consistent with the claim that none of the fishermen would be doing anything wrong by overfishing (in the absence of the established institutional order). But the intuition that there must be something

⁸¹ See, e.g., Fanciullo (2021), Pinkert (2015), Portmore (2018).

⁸² Portmore (2018: 329–337) offers good reasons to adopt a different, non-standard formulation of the principle, but discussion of which formulation we should accept would take us too far from the main object of the chapter. For my purposes it is sufficient that it becomes at least somewhat intuitive why we want to find some fault with what agents in collective harm cases do.

wrong about acting in ways that bring about the harm to a large extent motivates the search for a solution.⁸³

Nefsky (2015) endorses a stronger version of the restriction advanced in the previous paragraph. She claims that the wrongness of collectively harmful acts must relate to the fact that these acts are *instrumental in producing the harmful outcome*. She argues that unless we show that an act is not superfluous in contributing to the harmful outcome, we would not be able to adequately identify a reason not to perform this act. So, in her view, it would not be enough to point out that using a big net too often is wrong because this way a fisherman exceeds his fair share of fish or because he thereby becomes a part of a group that causes collective harm. I mention Nefsky's proposed restriction to make it explicit that I shall not assume it. I only endorse the weaker restriction mentioned in the previous paragraph, namely, that a solution should identify the wrongness in *acting* in collectively harmful ways. But I do not make the stronger claim that acting in these ways is wrong *because these acts are instrumental in bringing about the harm*.

4.3 Kantian Imperfect Duties as a Basis for a Solution

One interesting proposal holds that agents have an imperfect duty not to act in collectively harmful ways. What this means depends on one's account of imperfect duties. I shall consider a broadly Kantian account of imperfect duties that I have outlined in Chapter 1. There are three main reasons for paying attention to this account. First, it seems to do a good job in satisfying both the desiderata. Second, the account of Kantian imperfect duties that I consider is widely accepted, and so it would be a nice result if this conception also helps in solving the inefficacy problem (I think it does not). Third, Albertzart (2019) proposes a solution to the inefficacy problem that relies on this account.

⁸³ Again, if the reader does not share these intuitions, this would not be a problem for the soundness of my main argument. It would just make it somewhat less interesting.

Here is a very brief reminder of the Kantian end-based view. Imperfect duties are duties to adopt and promote certain ends. These duties directly require cultivating or forming certain attitudes that are constitutive of the relevant ends and indirectly require acting in accordance with these attitudes. A paradigm example of an imperfect duty is the duty to adopt and promote the end of others' happiness. If failing to perform an act is inconsistent with having one or more of the relevant attitudes, then performing this act is strictly required. When an agent has an option of performing such an act, I shall say that this agent has a 'golden opportunity' to promote one of her obligatory ends (Noggle 2009). Perfect duties prescribe or proscribe performance of acts of certain types. Classic examples of perfect duties are the duty to keep one's promises and the duty against harming others. I should make it explicit that in this chapter I consider only the imperfect duties related to the obligatory end of other people's happiness or welfare as this obligatory end seems most relevant to cases of collective harm. At least intuitively, if we are to locate the failure to adopt an obligatory end in acting in collectively harmful ways, the relevant end must be somehow related to others' welfare.

In a moment, I shall explain how this conception can serve as a basis for a solution while satisfying the desiderata. But first, to address the elephant in the room: Why are imperfect duties relevant in collective *harm* cases at all, given that the duty not to harm is paradigmatically perfect and the imperfect duties I discuss are usually thought to be concerned with promoting others' welfare? If it were clear that individual agents violate their duty not to harm, the problem would not require much discussion. But it has been notoriously difficult to argue that this is so. It is difficult to argue that agents are violating the duty not to harm precisely because it is difficult to show that any individual agent contributes to the collectively caused harm. At best, individuals have a small chance of making a difference, and even that is not obviously true. So, since it is very difficult to claim that individual agents in collective harm cases are harming anyone, or putting others under non-negligible risk of harm, it is difficult to argue that they are violating their duty not to harm. Furthermore, it would be an oversimplification to claim that imperfect duties that are concerned with others'

welfare amount to requirements simply to promote others' welfare. They are more complex than that, which is another reason to consider them seriously as a basis for a solution.

Let us now see how a solution based on this view on imperfect duties would satisfy the desiderata for resolving the problem of inefficacy, and what such a solution might look like. A solution, to even count as such, must identify the moral wrongness in something that individual agents do when they collectively cause harm. In the case of the relevant imperfect duty, the wrongness would be in failing to adopt the obligatory end of others' happiness or in failing to act in a way consistent with having the attitudes constitutive of this end. So, to satisfy the desiderata, the imperfect-duty-based solution must show that by acting in collectively harmful ways, individual agents are doing something that is inconsistent with adopting the relevant attitudes. Let's consider each desideratum in turn.

Recall that the first desideratum is that a solution must *in some cases* classify individual acts of the relevant kind as pro tanto wrong. So, for the imperfect-duty-based solution to satisfy this, it must sometimes judge specific kinds of acts to be inconsistent with having the relevant attitudes. And there seems to be a clear way in which it could do that. It is an important feature of the Kantian account that in some cases, golden opportunities arise. There are cases in which acting in a specific way is strictly required, because failing to act this way would be inconsistent with having the relevant attitudes. And abstaining from performing certain collectively harmful acts could count as taking advantage of the golden opportunity. Let's consider the case of one-off elections again. It could be argued that voting is a golden opportunity to act consistently with the obligatory end of others' happiness. This is so even if it is assumed that one's vote is extremely unlikely to make a morally significant difference. Several reasons are relevant for concluding that failing to vote could be inconsistent with adopting the obligatory end. One reason is that electing the worse candidate would be a result of insufficiently many votes cast for the better one and each voter knows that. Another reason is that we can assume that voters have no strong moral justification for failing to vote. Finally, each voter can only vote once, and each voter

knows that. It then could be argued that failing to vote, voters would exhibit a lack of proper concern for others' happiness. Knowing that the failure to get enough votes is directly causally relevant to bringing about the harmful outcome, that there is no good justification to abstain from voting, and that each of them can only vote once seems sufficient to conclude that failing to vote is inconsistent with having a proper concern for others' happiness. This might not seem too convincing, but my aim is not to defend the view that performing a single act of the relevant kind does sometimes entail a violation of an imperfect duty. My aim is to show that an argument along these lines could be constructed and that there is some plausibility to it. So, it is enough for my purposes that there is an intuitive case for thinking that an imperfect-duty-based solution can satisfy the first desideratum.

Let's consider whether and how such a solution can satisfy the second desideratum. Recall that this desideratum holds that acts of the relevant kinds must not be classified as pro tanto wrong in *all* cases. So, there sometimes must be other ways to condemn acting in collectively harmful ways. For instance, it could be that while individual acts of the relevant kinds are not pro tanto wrong, patterns of activity involving several such acts are pro tanto or all-things-considered wrong. The imperfect-duty-based solution seems to be well-suited to condemn patterns of activity. Let's consider *Fishermen* again. Let's assume that if each fisherman uses a big net just once, the harmful outcome does not occur.⁸⁴ The fish population will only decline substantially (leading to famine) if enough fishermen use big nets enough times. On this assumption, it is difficult to see how the fact that one fisherman uses a big net just once would be inconsistent with his taking others' happiness seriously. This is so even if he uses his big net without a strong moral justification, but for a less weighty reason (for example, that he wants to make some extra cash from selling fish to buy his child a nice birthday present).

⁸⁴ Apart from the harm to the fish they catch, but we're abstracting from this consideration for the present purposes.

However, if we consider a fisherman who uses a big net regularly without having a weighty moral reason to do so, it becomes much easier to argue that he fails to take others' welfare seriously enough and thus fails to adopt it as an end. Engaging in a collectively harmful activity regularly, without any restraint, and without a strong moral justification does seem to imply a lack of concern for others' welfare. Note that it could be argued that the fisherman exhibits a lack of concern for others by using a big net too often even if his pattern of activity is still quite unlikely to make any difference to the harmful outcome. The mere fact that acts of this type could sum up to bring about this outcome, and the fact that he does not seem to care about this seem sufficient to yield the conclusion that he does not have others' happiness as an end. So, it seems that the imperfect-duty-based solution could condemn the actions of those fishermen who use big nets too often.

It seems then that the imperfect-duty-based solution can satisfy both desiderata. It seems to be able to classify acts of certain kinds as pro tanto wrong in some cases, while only condemning patterns of activity in other cases (without condemning individual acts). Moreover, it respects the implicit restriction that an acceptable solution must find wrongness in the *acts* of the relevant kinds. It does not, for instance, imply that the fishermen violate their obligation to come up with an institutional arrangement that would minimise the risk of depleting the fish population, but do nothing wrong apart from that. It does not imply that they do nothing wrong, for instance, if they keep on using big nets without restricting themselves in the absence of such an institutional arrangement.

So, it appears that a proponent of the imperfect-duty-based solution has an intuitive case in its favour. However, in the next section, I shall argue that there are significant problems with a solution like this.

4.4 Why Kantian Imperfect Duties Are Unfit as a Basis for a Solution

I do not think that acting in collectively harmful ways will often amount to a violation of an imperfect duty. My strategy to argue for this is as follows. I shall argue that it is impossible to make any judgments about whether agents violate their imperfect duties when we consider only isolated patterns of their activity, as we often do in collective harm cases. I then show that if we extend the scope of our assessment across many such patterns, it turns out that acting in collectively harmful ways often has no impact on whether an agent violates her imperfect duties.

4.4.1 Obligatory Ends and Isolated Patterns of Activity: The Problem of Scope of Assessment

Collective harm cases, as they are discussed in the literature, often take the form of carefully crafted scenarios. Many important papers that wrestle with the inefficacy problem, attempt to offer a general solution that extends to all cases of the relevant kind by considering various isolated cases that often take the form of thought experiments.⁸⁵ What is common about all these discussions is that they attempt to figure out a solution by investigating one representative case in isolation, while assuming that the solution would generalise to all other structurally similar cases. But we cannot draw any conclusions regarding agents' success or failure to adopt the obligatory end by considering one isolated case. Why not?

Consider *Fishermen* again. In the previous section, I have suggested that a proponent of an imperfect-duty-based solution could argue that if one of the fishermen, call him Frank, often uses a big net without a good moral justification, this

⁸⁵ Most common ones are Parfit's (1984) *Harmless Torturers* and *Drops of Water*. For discussions of Parfit's cases see Arntzenius & McCarthy (1997), Fanciullo (2021), Nefsky (2017), Otsuka (1991), Spiekermann (2014). See also Albertzart (2019), Kagan (2011), Tiefensee (2022).

entails that Frank fails to have adopted the end of others' happiness. Recall, however, that whether one has adopted the relevant end depends on one's having formed attitudes, such as the resolution to treat others' happiness or well-being as a source of reasons and an intention to promote it. And one's behaviour in one domain of activity is not sufficient to determine whether one has the relevant attitudes. Patterns of behaviour must be assessed across domains of activity. In Frank's case, using a big net is just one such pattern. But he also acts in many other ways. There are facts about how he treats his family and friends, his co-workers, his acquaintances, and strangers too. And these facts are directly relevant for assessing whether he adopts the relevant ends and whether he does enough to promote it. The fact that he often uses a big net is insufficient to conclude that he fails to adopt the relevant attitudes. More generally, the key point is that the fact that one's actions do not seem to exhibit the right kind of concern for others' happiness in one domain of activity *does not entail that one does not have the right kind of concern*.⁸⁶ Here is a case that may, by analogy, help drive this point home:

Monday Exercise. Peter claims that he wants to lead a healthy lifestyle. However, he consistently fails to attend his Monday morning sessions at the gym. He has no justification for doing that other than that he likes sleeping in on Monday mornings. Could it be concluded that Peter does not have the aim of leading a healthy lifestyle?

I don't think that such a conclusion is warranted. To make an assessment of whether Peter adopts such a lifestyle, it is insufficient to consider just one pattern of activity that he exhibits. We must consider many relevant patterns. We need to know whether he sticks to his diet, whether he exercises on other days, whether he tries to

⁸⁶ And even that might not always be enough to draw reliable conclusions. It is at least a logical possibility that an agent never has any opportunities to help others or to in any way exhibit the relevant attitudes, which makes it even harder to judge whether she has adopted the relevant end.

maintain a healthy sleep schedule, etc. Peter might be failing in one respect, but still be committed to his aim overall. Or, given that Peter succeeds in his overall goal, we might judge that he does not exhibit any irrationality by failing to attend his Monday sessions.⁸⁷ And the case could be flipped: he could be excelling in one isolated aspect while failing to count as having adopted the aim in virtue of failing in all other contexts relevant to his aim.

In sum, there is a problem with the scope of assessment. When evaluating whether an agent complies with her imperfect duty to adopt others' happiness as an end, it is insufficient to consider how she acts in a given collective harm case. We need to employ a wider evaluative scope and look at how an agent acts with respect to attending to others' happiness across a much wider pattern of behaviour.

4.4.2 *Why Imperfect Duties Do Not Help to Solve the Inefficacy Problem*

So far, I have identified the problem of the scope of assessment. We cannot judge whether an agent violates her imperfect duty just by looking at her actions in a single collective harm case. This makes the imperfect-duty-based solution more difficult to apply: after all, we need to consider a lot of facts about an agent to judge whether she complies with her imperfect duties. Considering her actions in an isolated collective harm case without taking those other facts into account is not enough. It is not, however, an insurmountable problem since it does not by itself entail that the imperfect-duty-based solution does not work.⁸⁸

⁸⁷ See Tenenbaum & Raffman (2012).

⁸⁸ Note also that even if we could make judgments regarding whether an agent adopts the end of others' happiness on the basis of her actions in one collective harm case taken in isolation, this would not guarantee that the solution works. It could be, for instance, that there is a mere empirical correlation between violating one's imperfect duty and acting in collectively harmful ways. This correlation would make it possible to claim that those who act in collectively harmful cases violate their imperfect duties. But for the success of the imperfect-duty-based solution to the inefficacy problem such a correlation is

To assess whether it works, we need to answer two questions. The first is whether an agent's actions in collective harm cases are at all relevant for determining whether she complies with the requirements of imperfect duties. If they are not, then the solution does not even get off the ground. This is because in such a case, her actions in collective harm cases would not in any way affect her success in complying with her imperfect duties. I believe we should assume that one's actions in collective harm cases are relevant for determining whether she complies with imperfect duties. Suppose that Frank the fisherman not only uses a big net without restraint, but also shows no restraint whenever he uses other kind of resources from limited pools, leads an emissions-heavy lifestyle, buys products made on factory farms, and acts in collectively harmful ways in many other structurally analogous cases. It is hard to believe that, *taken together*, his patterns of behaviour across these cases are completely irrelevant for determining whether he takes others' happiness seriously. So, let's assume that they are relevant.

However, answering the first question is not sufficient for the success of the imperfect-duty-based solution. Even if one's actions in collective harm cases are relevant, it does not follow that an agent's action in a single collective harm case can make a difference to whether this agent adopts the end of others' happiness. And for the success of the solution, it must be possible for an agent's actions in a single collective harm case to make such a difference. Consider the following pair of cases to see why.

Frank — Big Net (FBN). Frank the fisherman (from *Fishermen*) lives an ordinary life. He cares about his close ones to a reasonable extent, and he helps strangers on some occasions, but he does not actively seek out opportunities to help, nor does he help whenever doing so requires

not enough. It must also be the case that the reason why the relevant imperfect duty is violated lies in one's acting in collectively harmful ways.

bearing more than moderate costs. He also regularly uses a big net without a strong moral justification.

Frank — Small Net (FSN). All the details are just as in FBN, but Frank never or very rarely uses a big net without a strong moral justification.

Presumably, *almost* all the details that are relevant for determining whether Frank complies with his duty to adopt the end of others' happiness are held fixed across FBN and FSN. The only important difference between these two cases is that in the former, Frank frequently acts in an isolated collective harm case (by using a big net), whereas in the latter he does not. In line with the desiderata discussed in section 2, an acceptable solution to the inefficacy problem should condemn Frank's regular use of a big net in FBN but not his (rare or non-existent) use of a big net in FSN. So, the imperfect-duty-based solution should condemn Frank's use of a big net in the FBN but not in FSN on the grounds that he fails to adopt the end of others' happiness in the former but not in the latter case. But all the details, apart from his use of a big net, that are relevant for determining whether Franks complies with this imperfect duty are held fixed across both cases. Thus, if Frank violates his imperfect duty in FBN but not in FSN, it must be because of the difference in his use of a big net. So, whether the imperfect-duty-based solution successfully condemns Frank's actions in FBN but not in FSN depends on whether his pattern of using a big net in FBN could be sufficient to determine that he violates his relevant imperfect duty. In other words, it must be possible that his actions in an isolated collective harm case (*Fishermen*) make a difference to whether he adopts the end of others' happiness.

Here is the second question for assessing the success of the imperfect-duty-based solution: Can an agent's actions in a single collective harm case make a difference to whether one adopts the end of others' happiness? It is not obvious that they can. I find it plausible that even though one's actions across collective harm cases in which one participates are relevant for whether one adopts the end of others' happiness, a single pattern of behaviour in one specific collective harm case can make no difference to

whether one adopts or fails to adopt the obligatory end. After all, in the lives of most agents there is an abundance of factors that contribute to determining whether she has the relevant end. We can draw an interesting parallel here. The insensitivity of harmful outcomes to individual collectively harmful acts encountered within one collective harm case seems roughly analogous to the current discussion. Just as instances of using a big net, taken together, are relevant for determining whether a harmful outcome occurs, various patterns of behaviour across many cases, taken together, are relevant for determining whether an agent adopts the obligatory end. And just as an individual instance of using a big net can make no morally significant difference to the harmful outcome, a pattern of behaviour in one specific collective harm case can make no difference to whether an agent adopts the obligatory end.

If we accept that an agent's pattern of behaviour in a specific collective harm case cannot make a difference to whether this agent adopts the obligatory end, then the imperfect-duty-based solution does not work. It would never be able to condemn such a pattern of behaviour *on the grounds that* this pattern violated an imperfect duty. And so, we would have to conclude that agents in collective impact cases do nothing wrong even if they act in collectively harmful ways without restraint (or at least nothing that is wrong due to violating imperfect duties).

4.4.3 *Two Reasons Why Making a Difference Wouldn't Save The Imperfect-Duty-Based Solution*

For the sake of the argument, let's grant that a single pattern of behaviour can make a difference to whether one adopts the relevant end. So, let's grant that Frank violates his imperfect duty in FBN but not in FSN precisely because of his increased use of a big net in the former case. Does this vindicate the imperfect-duty-based solution? It does not. There are two reasons for this. Let's consider each in turn.

First, if a single pattern of behaviour can make the relevant difference, *why* can it make it? Why think that Frank's use of a big net in FBN is significant enough to determine that he does not have the relevant attitudes? One potential answer is that it is significant because it causally contributes to overfishing and thus to causing famine.

But this answer would make the imperfect-duty-based solution much less interesting. Once we assume that each act of the relevant kind in any collective harm case has a chance of making a causal contribution, we can condemn such acts on the grounds that they cause harm or at least have a chance of causing harm (and it is even easier to condemn patterns of such acts given that the chance of their causing harm is higher). It seems that we do not need extra considerations such as that performing these acts lead to violations of imperfect duties.⁸⁹

How can we explain why Frank's use of a big net amounts to a violation of his imperfect duty if we do not assume that his actions causally contribute to harm? Generally, it seems, that unless there is *some* moral reason not to use a big net (or act in any other collectively harmful ways), performing such acts is consistent with having the attitudes constitutive of the end of others' happiness. After all, why would acts against which there is no moral objection be inconsistent with having others' happiness as an end? Why would such acts exhibit a lack of concern for others' happiness? It seems that the assumption that there is some independent moral reason not to perform these acts is necessary for these acts to count as violating the imperfect duty to adopt others' happiness as an end. But if to conclude that these acts violate an imperfect duty we need an independent moral reason not to perform them, then the imperfect-duty-based solution to the inefficacy problem is superfluous or parasitic. Once we have those independent moral reasons, we already have a solution: acting in collective harm cases would be *pro tanto* wrong because it would go against those independent moral reasons.

But let's assume, for the sake of the argument, that the imperfect-duty-based solution is not superfluous. Suppose that an agent's acting in collectively harmful ways does make a difference to whether this agent adopts the obligatory end without there being an independent moral reason not to perform these actions. There is a second

⁸⁹ Note that this moral reason does not have to be related to making a difference to harm. It could be, for instance, that the reason is that these acts non-superfluously contribute to harm without making a difference (Nefsky 2017).

reason why the proposed solution does not work. The reason is that even though an agent's pattern of behaviour in a specific collective harm case *can* make a difference to whether this agent adopts the obligatory end, it does not follow that such a pattern *does* make the relevant difference. And it is the latter claim that is needed for the solution to succeed. Let's consider FBN and FSN once again. We may grant that it is Frank's use of a big net that provides a reason why he violates his imperfect duty. Then the imperfect-duty-based solution works when applied to a pair of cases like FBN and FSN. But how often do cases such as FBN and FSN arise? It seems to me that in the vast majority of collective harm cases an agent's pattern of behaviour in any one such case does not make a difference to whether the agent adopts the obligatory end. In most cases there are other important details that make a difference to whether the agent adopts the relevant end.

Suppose that Frank neglects all available opportunities to help others and does not show even a slight interest in others' well-being. In such a case, it is quite unlikely that his pattern of using a big net makes any difference to whether he adopts the obligatory end: his failure to adopt this end seems overdetermined. Conversely, if Frank lives his life selflessly, and spends a substantial part of his resources to alleviate others' need, then his using a big net very often is still plausibly consistent with his having adopted the end of others' happiness.

So, regardless of whether an agent violates her imperfect duty or not, this would most likely not be due to her actions in a single collective harm case. But if her violation of her imperfect duty is not due to her actions in a collective harm case, then the imperfect-duty-based solution to the inefficacy problem does not work. Whenever we would consider her actions in a collective harm case, they would almost never be wrong (at least not because they would amount to her violating the duty to adopt others' happiness as an end). The imperfect-duty-based solution is close to what has been classified in the Introduction to this chapter as a 'sceptical' response to the inefficacy problem.

4.4.4 *What We Are Left With*

Let's summarise the discussion so far. I have argued that we cannot base a solution to the inefficacy problem on an appeal to imperfect duties. I have argued that if we consider an agent's isolated pattern of activity in a given collective harm case, we cannot make any conclusions regarding her compliance with the requirements of imperfect duties. Once we extend the scope of the assessment to include many of this agent's patterns of activity, we see that two questions should be answered to assess the success of the solution.

One question is whether actions in collective harm cases are at all relevant for determining whether an agent adopts the end of others' happiness. I have assumed that they are relevant. The second question is whether these actions can make such a difference. It is not obvious that they can. And if they cannot, the solution fails to yield any non-sceptical verdicts about the wrongness of acting in collectively harmful ways.

I then argued that even if an individual's actions in collective harm cases can make such a difference, there are two reasons to think that the solution is still unsuccessful. One reason is that it is possible that these actions can make a difference to whether an agent adopts the obligatory end only if there is an independent moral reason to avoid performing them. But if there is such an independent reason, then it is already sufficient for solving the inefficacy problem, and the appeal to imperfect duties is superfluous. The other reason is that even if we assume that the appeal to imperfect duties is not superfluous, there are grounds for thinking that an agent's actions in an isolated collective harm case would often make no difference to whether he adopts the obligatory end. So, at best, the imperfect-duty-based solution works in a very limited number of cases, often yielding sceptical verdicts. And at worst, it never provides a non-sceptical verdict, thus failing as a solution to the inefficacy problem.

The result of this discussion is not completely negative. Failure to solve the inefficacy problem is not an objection to the account of imperfect duties that I have been discussing. Rather, it is a case of a mismatch between a problem and a potential solution. Imperfect duties are useful in shaping our moral deliberation when we

consider our lives over time and engage in planning and structuring our activities. The inefficacy problem is a theoretical puzzle that seems to be best addressed by considering isolated cases where certain relevant variables are isolated, which prevents us from saying anything useful about imperfect duties in the context of such cases.

To conclude the discussion, I would like to suggest that imperfect duties are relevant in addressing the *practical* side of the inefficacy problem even though they do not work in addressing the *theoretical* side. In other words, there is a reason to think that even though imperfect duties do not reliably condemn individual agent's actions in collective harm cases, they nevertheless can guide action in ways that would plausibly lead to avoiding the harmful outcomes. Let's start with the observation that in many real-world collective harm cases agents have the option of communicating with one another. In *Fishermen*, for instance, fishermen live in the same town, so it is very plausible that they can communicate. They are also aware that each of them has reasons to use a big net and that if everyone does that regularly, they will all suffer. Given this, it seems that if each of them takes others' happiness seriously as an end, he should take reasonable steps to communicate and cooperate with others to avoid the harmful outcome. So, their imperfect duties seem to direct them to establish an institutional arrangement that would prevent the harmful outcome from occurring. For instance, they might establish enforceable limits on the use of big nets or prohibit their use altogether. In this way, imperfect duties help in addressing the practical side of the inefficacy problem, namely, avoiding the harmful outcome. But this, admittedly, does not address the theoretical side. After all, if the fishermen use their big nets often in the absence of an institutional arrangement, this will not reliably constitute violating an imperfect duty.

Given that imperfect duties might seem to direct us to cooperation with others in collective harm cases, which seems to be what we need to do (inefficacy problem aside), one might think that this shows that the inefficacy problem is not even that serious a problem after all. Why isn't it that serious? The reason is that even though we lack the solution to the theoretical puzzle, the fact that a harmful outcome (such as

famine in *Fishermen*) occurs, suggest that at least some agents violated their moral obligations. These are the obligations to cooperate and come up with an institutional order that would help prevent harmful outcomes. This can also be argued to show that the moral harmony principle that in part motivated the need for a solution to the inefficacy problem is satisfied. It is satisfied because when collectively caused harmful outcomes occur, we can't say that everyone behaved as they should have: some people failed to take steps to communicate and cooperate. And if our imperfect duties direct us to avoid seriously harmful outcomes, this may seem sufficient. One might think that we should mainly care about avoiding harmful outcomes, which imperfect duties seem to help us do. Whether we can solve the theoretical puzzle is a separate, less pressing question.⁹⁰

While I do believe that avoiding disastrous outcomes is generally more important than solving theoretical puzzles, I am not convinced that we should abandon our attempts to solve the inefficacy problem. After all, it has not been established beyond any question that imperfect duties do direct us to trying to establish institutional arrangements that prevent harmful outcomes. This is a task for further work. But even if this were established, there would still be some reason to solve the inefficacy problem. Admitting that the problem persists does not force us to conclude that imperfect duties are not a useful element in moral theorising, only that they do not help us solve this particular problem.

4.5 Alibert's Proposal

In this and the next sections, I consider and argue against a different attempt to use Kantian imperfect duties to formulate a solution to the inefficacy problem. This alternative proposal can be seen as an attempt to argue that an agent's actions in an isolated collective harm case can and often do make a difference to whether an agent

⁹⁰ I say more about the theoretical puzzle in the next chapter.

adopts others' happiness as an end. This attempt merits separate discussion, since the reasons to reject it are independent from the ones mentioned so far.

Maïke Albertzart (2019) claims that the inefficacy problem can be solved by an appeal to Kantian imperfect duties. The general idea is that even if individual agents do not make a difference by acting in collectively harmful ways, they might be acting wrongly because they might be violating their imperfect duties. Her (2019: 7) construal of imperfect duties does not significantly differ from the one presented in section 4.3: "Whereas perfect duties prescribe the performance or omission of certain act types [...], imperfect duties are duties to adopt certain obligatory ends". Following Kant, Albertzart accepts that the two obligatory ends are other people's happiness and one's own perfection. Her understanding of obligatory ends and what it means to adopt them also seems sufficiently similar to the one presented above. Adopting an end is a matter of forming and cultivating certain attitudes, including forming an intention to promote the relevant end. So, indirectly, the requirement to adopt an end is a requirement to act in appropriate ways: one cannot be judged to have adopted the end of others' happiness if one never acts on one's intention to promote their happiness at least on some favourable opportunities.

Before I present the specifics of Albertzart's proposal, there are two conceptual clarifications to be made. First, Albertzart makes use of the notion of the necessary means to an obligatory end. She (2019: 8) claims, for instance, "To adopt the happiness of others as an end implies willing the necessary means for achieving this end". Talk of *achieving* the end of others' happiness is confusing as there isn't a clear sense in which this end can be achieved, unlike, for instance, an end of writing a book or leading a healthy lifestyle. This difficulty can be sidestepped. It is much clearer to think of what might be necessary to count as having adopted the end of other's happiness: it is necessary to form and cultivate certain attitudes like the intention to promote others' welfare and to act accordingly at least on some favourable opportunities. And since adopting an end implies wholeheartedly trying to achieve it, we should understand the necessary means to an obligatory end as attitudes that an agent must form and actions that she must perform to count as having adopted this

end. If an agent does what is necessary to adopt the end, she thereby does what is necessary to wholeheartedly try to achieve it, whatever achieving it means in the present context.

Second, Albrecht (2019: 8) sometimes speaks of *chosen means* to an obligatory end. By *chosen means*, she seems to understand what an agent actually chooses or commits to do as part of adopting or promoting her obligatory end. So, when an agent forms a general intention or resolution to help others and chooses a policy regarding how to act on this resolution, she chooses a means. There is at least one pattern of activity an agent must exhibit to count as having adopted an obligatory end, and there are often many possible patterns from which an agent may choose. It is necessary, if one wants to count as having adopted the obligatory end, to choose some such pattern. Once an agent commits to a certain pattern, it becomes her *chosen means* to this end. For instance, it is plausible that there are many possible policies regarding how to act on one's resolution to promote others' happiness. Choosing some policy among the available ones is part of the necessary means to the end of others' happiness, but only the policy that an agent actually chooses is part of her *chosen means*.

Having discussed the preliminaries, let us proceed to the core of Albrecht's proposal. Her (2019: 8) main claim is that a collectively harmful act is wrong if the universalised form of the maxim underlying this act contradicts the agents' *chosen means* to an obligatory end. Let us unpack this claim. Albrecht (2019: 5) understands a maxim as the subjective principle of an action that has the form "Do X in circumstances C for the reason R to reach the end E". To universalise a maxim is to imagine a situation in which everyone acted on this maxim, that is, performed the relevant act in the relevant circumstances for the relevant reasons, to achieve the relevant end. Albrecht is explicit about the fact that the contradiction test she proposes to solve the inefficacy problem is different from the familiar Kantian tests.⁹¹ There are two questions to be settled to fully understand and assess Albrecht's main

⁹¹ Here I am talking about the contradiction in conception and contradiction in willing tests.

claim. What would a contradiction of the relevant sort look like? And why would such a contradiction imply that this act is wrong?

It is helpful to address the first question with an illustration, which Albrecht provides (2019: 7–8). She considers the case of climate change, claiming that it is plausible that collectively harmful acts such as driving are subsumed under the obligatory end of other people’s happiness. In other words, avoiding such acts is part of promoting this obligatory end. So, while driving is not *pro tanto* wrong, it could be wrong, if acting in this way involved the relevant contradiction. An instance of driving could be wrong if the maxim underlying this act would contradict the means that an agent chooses to the end of others’ happiness. Albrecht argues that such a contradiction could arise. She first claims that combating climate change is a necessary means to other people’s happiness: one could not count as having adopted this obligatory end if one did not in any way resolve to combat climate change. So, an agent who wants to adopt the end of other happiness must resolve to combat climate change and to choose some means of combating it. Now, climate change could plausibly be combated at least to some extent by reducing unnecessary driving. So, we could assume that our agent wills or chooses that everybody avoid unnecessary driving as a means to combating climate change.⁹² If we further assume that this agent drives her car and does this on a maxim “I shall drive my car to reach my destination whenever it is convenient for me”, then this agent arrives at a contradiction of the relevant kind. As Albrecht (2019: 8) puts it, “[The agent] faces a contradiction between a chosen means for an obligatory end — namely, that everybody avoids unnecessary car trips — and the universalised form of her maxim, namely, that everyone drives their cars in order to reach their destination conveniently...”

Let us now move to the second question. Why does this contradiction imply that the act is wrong? Albrecht offers two considerations to answer this question. She (2019: 8) puts one of these considerations as follows: “...this contradiction expresses

⁹² As I argue below, this move is problematic – it is implausible that an agent could unilaterally will or choose things completely beyond her control.

a very Kantian and familiar idea: We should not single ourselves out for special treatment in the cases at hand. The agent would single herself out for special if she willed that everybody avoid unnecessary car trips as a means to combating climate change and at the same time acted on a maxim of driving her car whenever it is convenient for her.” So, one reason why the contradiction implies that the relevant act is wrong is that it is unfair to will that others abstain from acting in collectively harmful ways, while at the same time choosing to act in this way. The other reason is that this contradiction implies that an agent fails to adopt the obligatory end. Suppose that the agent chooses a certain means for an obligatory end, but then acts on a maxim whose universalised form is inconsistent with this means. In such a case, the maxim, on which the agent acts, contains an end that is inconsistent with the obligatory end. In the case of our agent, the end of her comfort that is implicit in her maxim of driving whenever it is convenient for her is inconsistent with her combating climate change, which serves as a necessary means for the obligatory end of others’ happiness.

While Albrecht’s discussion focuses on the case of climate change, she seems to intend her solution to be general: in any collective harm case, an agent would act wrongly if faced with the relevant contradiction. Facing the contradiction implies that the act would be wrong in at least two ways described above.

4.6 Searching for the Contradiction

I shall now argue that Albrecht’s proposal faces a significant problem. The problem is that the contradictions of the relevant sort are either practically impossible (although logically possible) or somewhat likely to occur but still very rare. If the former is true, then the proposed solution clearly fails since it would fail to identify any ways to condemn collectively harmful acts. If the latter is true, then the solution is not general enough. It would fail to condemn collectively harmful acts in most cases.

Whether the contradictions are virtually impossible or rare depends on how a certain claim, which is implicit in Albrecht’s solution, is interpreted. The claim is that

at least one of the available means to the end of others' happiness must be relevantly connected to avoiding the harmful outcome. For example, in Albrecht's climate change case, the connection is clearly present. The necessary means to the obligatory end is combating climate change. Combating climate change is the way of preventing or merely contributing to preventing the harms of climate change.

What if there were no connection? If there were no connection, then regardless of whether one did anything to combat climate change, there would be no difference to one's success or failure in adopting the obligatory end. In such a case, regardless of what one's maxim were when one acted in collectively harmful ways, there would be no contradiction. For example, if one acted on the maxim of driving whenever that would be convenient, one could still be considered to have adopted the end of others' happiness. This is because if there is no connection between avoiding the harmful outcomes (such as the climate change) and the end of others' happiness, one could act in collectively harmful ways while successfully adopting the obligatory end. There could be a contradiction of the relevant kind if the agent's chosen means to the obligatory end were somehow inconsistent with the universalised form of the maxim underlying one's collectively harmful actions. But it is hard to see what means to the end of others' happiness an agent could choose such that it would both be (i) not in any way connected to preventing the harms of climate change and (ii) inconsistent with everyone's driving whenever that would be convenient. This also shows that while having an option to will the means that is connected to preventing the harmful outcome is necessary for the contradiction to be possible, for such a contradiction to *actually* arise, the agent must *actually* will this means, that is, take this option.

So, Albrecht's solution must include the assumption that at least one of the means to the obligatory end would be relevantly connected to preventing the harmful outcome. There are two ways in which this connection can be construed. Depending on how it is construed, the relevant contradictions are either practically impossible or possible but rare:

- (I) Taking the means to the obligatory end *prevents* the harmful outcome.

- (II) Taking the means to the obligatory end *merely contributes to preventing* the harmful outcome.

If the connection is understood in the sense (I), the relevant contradictions are virtually impossible. If in the sense (II), they are not merely logically possible but still rare. I shall now discuss each construal in more detail.

4.6.1 When Contradictions of the Relevant Sort Are Virtually Impossible

To say that a certain means to the obligatory end *prevents* the harmful outcome is to say that willing or choosing this means is causally responsible for making it the case that the harmful outcome does not occur. Albrecht (2019: 7–8) seems to have this in mind when she says that combating climate change is at least one of the necessary means to the end of others' happiness and that one way to combat climate change is by willing that everyone avoid unnecessary car trips. For, presumably, if everyone did avoid such car trips, at least those harms of climate change attributable to excessive driving, would not come about.

Understanding the connection as one of *preventing* is problematic. The problem lies in the fact that one cannot coherently unilaterally choose or will a measure that would prevent a harmful outcome that is caused collectively. To understand why this is problematic, let us first discuss the sense in which an agent “wills” an event in this context. There is an unproblematic sense in which an agent could will events that go beyond her control. This is the sense in which we use “will” when we think about Kant’s categorical imperative and apply his contradiction in willing test, which invites us to test whether we could coherently or rationally *will* or *choose* a certain imaginary scenario to take place. When we use “will” in this sense, we, as Parfit (2011: 285) aptly puts it, “...suppose or imagine that we have the power to *will*, or choose, that certain things be true. We are doing a *thought-experiment*, which involves comparing different possible states of the world...” If Albrecht intended this use of “will” when claiming that agents can unilaterally will means to their obligatory ends such that these means go beyond the agents’ control, there would not be a problem. But that is not the

way Alibertart uses “will” in this case. To will or choose a means to one’s end is not to conduct a thought-experiment where one imagines that one has a choice between different possible worlds. To will a means to one’s end is to make a practical commitment in the actual world, to form an intention. In Kantian terms it can be put as regarding oneself as the cause of the end in question (Korsgaard 1996: 94). This is the sense in which Alibertart uses “will”. Crucially, she cannot use “will” in the other sense. After all, she discusses the means that agents choose when making decisions regarding how to promote their obligatory ends. She then must be talking about their *practical commitments* rather than about their engaging in merely comparing states of the world in the context of a thought experiment. And her use of “will” is problematic in the present context. It seems incoherent to *unilaterally intend* an event that one has little to no control over bringing about, such as the event in which everyone abstains from taking unnecessary car trips.

Why would this problem make the contradictions that Alibertart’s solution relies upon virtually impossible? Above, I have argued that contradictions of the relevant sort arise only if the agent has a means to the obligatory end such that this means is relevantly connected to preventing the harmful outcome and this agent in fact wills this means. I have also just argued that unilaterally willing a measure that *prevents* the harmful outcome is incoherent given that such outcomes are caused collectively. So, for a contradiction to arise, an agent must in fact will the means that she cannot coherently will. But then the contradictions of the relevant kind are logically impossible and not merely unlikely to ever occur in the actual world. I shall now discuss a potential response by Alibertart, which defends the logical possibility of these contradictions, but fails to establish that they could ever occur.

Alibertart acknowledges that there is a problem with intending events that one has little control over. But she suggests that the distinction between intending-to and intending-that discussed by Michael Bratman (2014: Ch. 3) could help defend her argument. Specifically, she claims that although an agent cannot unilaterally intend *to* bring about an event that involves action on the part of other agents, this agent can unilaterally intend *that* something is done even if this action can only be performed

together with other agents. If Albrecht's response is successful, she can claim that an agent could coherently will *that* everyone avoid unnecessary driving and that the contradictions of the relevant sort are possible after all.

Albrecht is correct in claiming that an agent can unilaterally form an intention—that with respect to actions that she cannot perform by herself. For example, I can coherently intend that you and I go rowing together, even though I cannot execute the intended activity by myself: it requires action on your part. However, there are further conditions on my intending *that* we go rowing that must be met for such intending to be possible. I am going to discuss just one condition, which Bratman (2014: 64) calls the “settle condition”. Discussing just this condition is sufficient, because intentions that have preventing collectively harmful outcomes as their objects are overwhelmingly unlikely to ever satisfy it.

The settle condition states that for me to intend that we go rowing, I must believe that my intention really does settle whether we go rowing. This condition is plausible because, if we understand intentions or willings as commitments to act, for such attitudes to be coherent, they must have as their objects only issues that are in some sense up to the agent, that this agent can settle (Velleman 1997: 32; Bratman 2014: 64–65). But how could it be that I both know that performing the action is not fully up to me and yet believe that my intention settles whether it is performed? Bratman's explanation is that the intentions of each participant in a joint activity are interdependent in a certain way. Roughly, each participant will continue to have the relevant intention only if other participants continue to have it as well, and if the collective action is performed, it is performed in virtue of the relevant intentions. So, my intention that we go rowing leads to our going rowing in part by supporting your intention that we go rowing, and vice versa (Bratman 2014: 66). So, my intention settles whether we go rowing by supporting your intention that we go rowing. It is up to me whether we go rowing because it is up to me whether to support your intention that is necessary for us to go rowing together. And it is also up to you for the same reason: because your intention supports my intention that we go rowing.

Furthermore, that their intentions are interdependent in the relevant way must be known by each participant for them to be able to coherently hold intentions with collective activity as its object. And if any of the participants learned about the change in others' intentions, she would update her own intentions. For instance, if I learn that you are no longer intending that we go rowing, it would be incoherent for me to still intend that we do it.

Crucially, the settle condition is too demanding to make it coherent to intend actions that could prevent harmful outcomes in collective harm cases. Cases of collective harm involve large unstructured collectives, where the relevant interdependence of intentions is nearly impossible to occur. Take the case of the climate change, which plausibly involves populations of several countries, if not the whole world. It is very hard to believe that the object of the intention, whose content involves action on the part of millions of people, could satisfy the settle condition. It is plausible that even in cases of much smaller scale, like the *Fishermen* case, these conditions would likely fail to be met. This is not surprising, given that Bratman's theory is designed in the first place for cases of what he calls "modest sociality" (2014: 8), that is, small-scale shared agency.

We should conclude that even though it is logically possible for an agent to form intentions whose content includes preventing harmful outcomes in collective harmful cases, it is hard to believe that such intentions could ever be coherently formed. Since one could almost never intend that a harmful outcome be prevented, contradictions of the relevant sort could almost never arise, and they would definitely fail to arise in most important cases such as climate change.

At this point, Albrecht might try a different response. She could argue that intending that everyone abstain from unnecessary driving is similar to having an ambitious intention that world peace is achieved. And such intentions might seem coherent. When one wills or intends such outcomes as everyone's abstaining from unnecessary driving, one is expected to take steps towards them rather than secure them. For instance, one might be expected to "form an intention to convince others to

use public transport, to vote for a political party that supports her cause, or to campaign for institutional changes” (Alberzart 2019: 9).

However, it is incoherent to have an intention such that you could not possibly realise the outcome written in its object or such that you could at best make a modest contribution to realising it. It is incoherent to claim that you intend that everyone abstain from unnecessary driving, when you really intend only to try to convince others to do so, and when you could not realistically hope to convince everyone. It is, however, possible to have ambitious ends, which one might not have the power to achieve. It is incoherent to intend to achieve world peace, but not incoherent to have it as one’s life goal or aim. In the case of having such a goal, one could coherently intend less ambitious activities that would contribute to reaching this goal. But intending things like preventing climate change-related harms or achieving world peace is incoherent.

4.6.2 When Contradictions of the Relevant Sort Are Rare

Suppose now that the connection between the means and the harmful outcome is understood in the sense that taking the means to the obligatory end *merely contributes to preventing* the harmful outcome.

On this construal, the relevant contradictions could arise. For example, an agent could will to change her lifestyle to a less emission-intensive one as part of combating climate change. This would plausibly count as making at least a modest contribution to the reduction of harm. And if this agent at the same time acts on maxims such as “I shall drive whenever convenient”, she does seem to face a contradiction: her maxim is incompatible with her chosen means to combating climate change, which is, in turn, necessary to count as having adopted the end of others’ happiness. This contradiction would be even more straightforward than the one proposed by Alberzart because it arises even without universalising the agent’s maxim. If this contradiction shows that the agent does something wrong, then Alberzart’s solution can condemn individual agents acting in collectively harmful ways.

However, such contradictions would still be quite rare. The reason for this is that there are several kinds of cases in which they would not arise. Furthermore, cases of

these kinds, arguably, would be more common than the ones in which such contradictions would arise. In one kind of case, contradictions would not arise, because an agent could choose some means to the end of others' happiness such that this means is unconnected to the harmful outcome. An example of such a case is not hard to come by. Consider *Fishermen* again. An individual fisherman, as any other moral agent, has an imperfect duty to adopt the obligatory end of others' happiness. He could, as part of the means to this end, will to contribute to avoiding the harmful outcome of depleting the fish population. However, he need not do so — he need not make this his chosen means. He may adopt and promote the obligatory end without willing anything related to avoiding the harmful outcome. He could try to help others in other ways. He could, say, volunteer as a lifeguard in his free time or invest the profits from catching the fish into building a new hospital or a school for the town. The point is that he could act on a maxim that, if universally acted upon, would lead to a harmful outcome and fulfil his imperfect duties by choosing a means to others' happiness that is unrelated to the collective harm case in question. There would be no contradiction that is supposed to explain the wrongness of his act despite it being intuitively wrong for the fishermen to use big nets given the harm resulting from their doing so. While I have discussed just one case, I think that many collective harm cases are similar to it in the sense that they do not involve the contradiction that is needed to condemn individual collectively harmful acts.⁹³

Alberzart's own example of climate change can also be plausibly recast as an instance of cases of this kind. The contradiction between the universalised form of the motorist's maxim and her chosen means for the obligatory end rests on the assumption that combating climate change is necessary to count as having adopted the obligatory

⁹³ This line of objection to Alberzart's argument is broadly analogous to the line taken in section 4.4.3. Namely, even if an agent's actions in a collective harmful case could make a difference to whether one adopts the obligatory end, they are quite unlikely to make such a difference. Similarly, even though an agent could will the means to the obligatory end such that it is inconsistent with the universalised form of his maxim, it is quite likely that she would will the means that cannot generate the contradiction in virtue of having no connection to avoiding the harmful outcome.

end. However, this assumption is controversial. One could imagine people fighting famine, working towards eradicating malaria and other diseases, or engaging in some other activity that is highly beneficial for others, while leading an emissions-heavy lifestyle. It seems counterintuitive to say that such people fail to take the end of others' happiness seriously enough because, despite all the good they are doing for others, they do not actively combat climate change. So, it is plausible that the motorist might act on a maxim to drive whenever it is convenient for her without failing to adopt and promote the end of others' happiness.

There is another kind of case. It could be that even though an agent chooses a means that is connected to avoiding the harmful outcome, this means is such that the relevant contradiction does not arise. In *Fishermen*, a fisherman could, as one of his means to the obligatory end, will to contribute to avoiding the famine from overfishing. However, he might without contradiction use his big net without limiting its use, acting on the maxim of using a big net because it brings him profits. At the same time, he could act on his intention to contribute to avoiding the harmful outcome by investing his profits in an alternative food source for the town. Despite his acting in collectively harmful ways, Albrecht's approach would fail to condemn his actions.

If what I have said so far is correct, Albrecht's proposal does not offer a general solution to the inefficacy problem. I have offered two kinds of cases, in which, if we take the inefficacy problem seriously, agents' actions should be condemned, but in which the contradiction of the relevant sort does not arise. My conjecture is that cases of the sort that I have described are quite common. If they are, then Albrecht's solution systematically offers sceptical verdicts. If the cases of the kinds that I have described are rare, then while Albrecht's solution does not systematically offer sceptical verdicts, it still fails as a general solution to the inefficacy problem. And, crucially, it fails in important cases such as the case of climate change.

4.7 Conclusion

In this chapter, I have discussed whether Kantian imperfect duties could form the basis of a solution to the inefficacy problem. I have argued that they generally do not fare well in this respect, given that it is possible for agents to act in collectively harmful ways and thereby cause significantly harmful outcomes without violating their imperfect duties. This is so partly because we can't make judgments about whether agents comply with their imperfect duties when collective harm cases are considered in isolation, and partly because if these cases are not considered in isolation, agents are not required to restrict their acting in collectively harmful cases across all such cases in which they participate. I have suggested that this result is not to be regretted, as it does not by itself constitute an objection to Kantian imperfect duties. Moreover, there is a way to argue that our imperfect duties may direct us to cooperate and come up with institutional arrangements that would prevent the occurrence of harmful outcomes in collective impact cases. Some may think that this practical result is all we need from a 'solution' to the inefficacy problem. Still, it is important to realise that the theoretical puzzle persists and that Kantian imperfect duties do not help to solve it.

I have also considered Albrecht's version of a solution based on Kantian imperfect duties. Her key claim is that acting in collectively harmful ways is wrong when a universalised form of the maxim that underlies an act of this kind contradicts the means that an agent wills or chooses to one of her obligatory ends. I have argued against this solution. My strategy has been to argue that depending on how specifically Albrecht's proposal is understood, contradictions of the relevant sort are either logically possible but virtually never occur or that they occur in some cases but are still very rare. This implies that Albrecht's solution yields sceptical verdicts in many cases: it fails to condemn many intuitively wrong collectively harmful acts.

Chapter 5

There Are No (Purely Phenomenal) Collective Impact Cases in Which You Can't Make a Difference

5.1 Introduction

In the previous chapter, I discussed and discarded the imperfect-duty-based solution to the inefficacy problem that arises in collective impact cases. In this chapter, my aim is to argue that this is not as regrettable as it may at first seem. This is because there are different kinds of collective impact cases and, as I shall argue, a solution is not required for at least one kind of collective impact cases because they cannot arise (or are at least very unlikely to arise). I have so far largely ignored the distinctions between different kinds of collective impact cases, but these distinctions are relevant to my discussion in this chapter. I shall first introduce these distinctions, explain why some kinds of case are more problematic than others, and zero in on one specific kind of case. After this, I shall argue that there are good reasons, partly grounded in empirical evidence, to think that this kind of case cannot arise. If my argument is successful, then the inefficacy problem is at least somewhat less challenging than it may appear.

My project in this chapter is broadly in line with attempts to resolve the inefficacy problem that argue that it is incoherent to claim that individual actions never make a

difference.⁹⁴ However, my argument is different from existing arguments. The key distinctive feature is that it relies on a certain empirical claim about how our ability of perceptual discrimination works, which has not been made use of previously. Furthermore, my argument does not rely on the assumptions of these other arguments, and thus avoids some of the problems with them.⁹⁵

5.2 Kinds of Collective Impact Cases: Setting the Focus

First, let's distinguish between threshold (or triggering) and non-threshold (or non-triggering) collective impact cases. I briefly outlined this distinction in the previous chapter, but it would be good to discuss it in more detail. Threshold cases are thus called because, in such cases, a sufficient number of acts of the relevant type trigger a morally significant effect upon hitting a precise threshold or boundary. For instance, in a case of voting in a simple majority system, many votes have a chance of crossing the threshold that decides the results of the election, and one vote becomes the triggering act. This is, of course, on the assumption that the outcome is 'open', that is, the result of the election is not overdetermined. If it was overdetermined, no single vote could be argued to make any difference. Non-threshold cases are such that for any single act of the relevant type, there is no chance that this act by itself triggers a morally significant effect. A common example is the case of climate change. It is thought that different acts that involve emitting greenhouse gases collectively cause harmful effects such as droughts, but a single act of emitting (such as going for a short drive) might

⁹⁴ Here are some notable examples of attempts to argue this way: Arntzenius and McCarthy (1997), Barnett (2018), Kagan (2011), Otsuka (1991), Regan (2000), Voorhoeve and Binmore (2006), Chappell (*unpublished manuscript*).

⁹⁵ For instance, my argument is not a version of the popular "reports-based argument" advocated by Arntzenius and McCarthy (1997) and Voorhoeve and Binmore (2006). For the sake of brevity, I shall not offer a detailed review of this argument or explain precisely how my argument differ from it.

not even have a chance of being the threshold-crossing or triggering act, because the relevant threshold is not sharp, but vague. Any act that involves emitting some sufficiently small amount of greenhouse gases would either not be sufficient to trigger the harmful effect or it would be superfluous due to there already being more than enough emissions. Other common examples are cases popularised by Derek Parfit (1984), for example, the *Harmless Torturers* (I discuss a version of this case below).

It is often thought that threshold cases are less problematic than non-threshold cases, primarily because a calculation of expected utility could provide a solution to the inefficacy problem in many, if not all, threshold cases. Even if the probability of one's vote deciding the result of the election is very low, the potential benefit of doing so could be so great, that the expected utility calculation would recommend voting. This solution is unavailable for non-threshold cases. If an act has no chance of making a difference, it cannot be recommended by an expected utility calculation.⁹⁶

My focus in this chapter is on non-threshold cases. But I shall not argue that cases of this kind cannot arise. There is a second distinction to be drawn. There are 'purely phenomenal' and 'not purely phenomenal' varieties of both threshold and non-threshold cases.⁹⁷ The former are such that the only morally significant feature of potential outcomes are qualitative aspects of experience. The latter are such that the outcome has non-experiential morally significant features (but may also have experiential ones). A classic example of a purely phenomenal case is Parfit's (1984: 80) *Harmless Torturers*:

⁹⁶ Things are more complicated. Budolfson (2018; *unpublished manuscript*) and Nefsky (2021) argue that even in threshold cases the expected utility approach does not always help, and thus such cases still suffer from the inefficacy problem. Still, non-threshold cases are more problematic regardless of how well the expected utility approach works in threshold cases. Tiefensee (2022) convincingly argues that this is so because to offer a solution to the problem in non-threshold cases we would need to revise our commonly held views on causation and to develop an account of decision-making under indeterminacy (and not merely uncertainty, as in threshold cases).

⁹⁷ I draw the inspiration for the labels "purely phenomenal" and "not purely phenomenal" from Nefsky (2011), although she uses a slightly different terminology.

Harmless Torturers (HT). There are a thousand torturers and a thousand victims. Each of the victims is attached to a device that can apply an electrical shock to their bodies. There are a thousand and one settings on the device, each setting is denoted as $S_0, S_1, S_2, \dots, S_{1000}$. Each setting corresponds to a certain intensity of the electric current. There is a button on the machine which, if pressed, moves the setting up by one and increases the voltage by some tiny increment. At the beginning, the machine is switched off and is at the setting S_0 . At S_0 there is no electrical current and the victims feel no pain whatsoever. If the button is pressed a thousand times, then at S_{1000} the victims will feel excruciating pain because of the electrical current. The victims cannot feel the effect of a single pressing of the button, that is, there is no perceptible difference in pain experienced by the victims when they move from one setting to the adjacent one. Suppose that each torturer can press their button only once. Then, although each of the torturers produces an effect that cannot be felt by any of the victims, the effect of their acts combined inflicts severe pain on all the victims. There is no harm done to the victims apart from the pain inflicted with the device and there is nothing bad about the pain except for the way it feels.

Since the harm of painful experience is the only morally significant effect that results from the torturers' actions in HT, this is a purely phenomenal case. With some modifications, it could be turned into a not purely phenomenal one. For instance, if by pressing buttons, torturers not only increased the intensity of the electric current, but also caused cancerous tumours to grow in each victim, then the case would be a not

purely phenomenal one (assuming, of course, that having a cancerous growth is bad for a person even at the stages when this does not affect her experiences.)⁹⁸

My argument in this chapter focuses on non-threshold purely phenomenal collective impact cases (NTPCs). There is a reason to think that purely phenomenal are more challenging than not purely phenomenal cases. The reason is that in the former, while each act makes *some* difference (e.g., to the intensity of the current), no single act makes a morally significant difference since, by stipulation, the effects of a single act cannot be felt or experienced and the only morally significant aspects in these cases are qualitative aspects of experience. So, purely phenomenal cases serve as hallmarks of non-threshold cases. However, when one encounters a not purely phenomenal case, one may doubt whether it is genuinely a non-threshold case, unless one assumes so for the sake of the argument. For instance, it is not unreasonable to think that even a tiny increase in the size of one's cancerous growth constitutes a tiny detriment to one's health. If that were so, then it would be a threshold case: each incremental increase would cross one of the many thresholds, making the relevant person's well-being worse and worse. But this, of course, does not show that there are no genuine non-threshold not purely phenomenal cases. Whether or not there are such cases, my argument does not extend to them. The possibility of non-threshold not purely phenomenal cases limits the significance of my argument somewhat. But the argument remains significant. Establishing that we have good grounds to think that NTPCs are impossible is by itself a significant result.

My discussion in the remainder of this chapter is structured as follows. Section 5.3 demonstrates how non-threshold cases seem to generate a Sorites series and risk running into the Sorites paradox. Section 5.4 shows that there is a way to preserve the possibility of NTPCs while avoiding the paradoxical conclusion. This strategy rests on accepting that the "indistinguishable from" relation is nontransitive. In section 5.5, I present my case against NTPCs. I argue that there is a better alternative explanation of why phenomenal cases avoid the paradox and this explanation is not consistent with

⁹⁸ The cancerous growth example is also borrowed from Nefsky (2011).

the possibility of NTPCs. Part of the reason why this explanation is better is that it has empirical support. In section 5.6, drawing on Diana Raffman's (2012) work, I present this empirical evidence. Section 5.7 summarises the discussion and concludes with some thoughts on why my argument does not extend to non-threshold cases more generally.

5.3 Getting to the Paradox

A single collectively harmful act does not even have a chance of making a difference iff the morally relevant features of the case are insensitive to changes caused by this act. In HT, each pressing of the button fits this description: while a single pressing of the button causes the intensity of the current to go up, it makes no difference to the pain the victims experience. The insensitivity of the pain experienced or felt by the victims to changes caused by a single pressing of the button can be intuitively expressed in the following way:

Tolerance (T). For all n , if the victims experience no pain at S_n , then the victims experience no pain at S_{n+1} .

This principle shows that no matter how many times the button has been pressed, pressing it just one more time will not make a difference with respect to the pain experienced by the victims. T is an instance of a tolerance principle. Tolerance principles are thus called because they show that the applicability of a vague predicate is tolerant or insensitive to small changes (Sainsbury 1988: 28). Or, as in the present case, the morally relevant feature is insensitive to changes caused by a single act of the relevant type. And principles such as T are problematic because they, together with

other plausible premises, give rise to the Sorites paradox.⁹⁹ Here is how the paradox can be constructed with T:

- (1) Victims experience no pain at the setting S_0 .
- (2) For all n , if the victims experience no pain at S_n , then the victims experience no pain at S_{n+1} .
- (3) By (2), if the victims experience no pain at S_0 , then the victims experience no pain at S_1 .
- (4) By (1) and (3), the victims experience no pain at S_1 .
- (5) By (2), if the victims experience no pain at S_1 , then the victims experience no pain at S_2 .
- (6) By (4) and (5), the victims experience no pain at S_2 .
- ...
- (1004) By (2), if the victims experience no pain at S_{999} , then the victims experience no pain at S_{1000} .
- (1005) By (1004) and (1003), the victims experience no pain at S_{1000} .

Because of the insensitivity of pain perception to changes in the intensity of the current caused by a single pressing of the button we have to conclude that the victims feel no pain at the setting S_{1000} . This is a highly implausible result. It also contradicts the setup of the thought experiment, according to which the victims are in excruciating pain at S_{1000} .

So, non-triggering cases, and purely phenomenal cases specifically, seem to generate the Sorites paradox. This raises a problem for those who want to claim that these cases are coherent. Moreover, most prominent solutions to the paradox, such as indeterminist or epistemicist solutions, rely on rejecting the tolerance principle, which is expressed by premise (2) in the present case (Keefe 2003; Sainsbury 1988). But this

⁹⁹ See Barnett (2018: 5–6), Elson (2016: 487–488), Nefsky (2011: 383–384, 2019: 6), Parfit (1984: 78–82).

principle seems necessary for the possibility of such cases because it is precisely this principle that expresses the idea that the pain experienced by the victims is insensitive to changes caused by a single pressing of the button.

This is a serious problem for the defenders of NTPCs. However, I shall argue that it does not follow that the soritical argument should be treated as a reductio and that we should conclude that purely phenomenal cases are incoherent.¹⁰⁰ This is because, as I shall argue in the next section, T is not needed to account for the possibility of NTPCs. A weaker principle would suffice.

5.4 Defending the NTPCs While Avoiding the Paradox

I shall now argue that T is not necessary to capture the insensitivity necessary for the possibility of NTPCs. Hence, a defender of NTPCs could argue that asserting the possibility of such cases does not lead to paradoxical conclusions.

The general proposal rests on weakening T, which is instantiated by premise (2) of the soritical argument.¹⁰¹ Weakening this principle involves three steps. The first step is to make explicit a certain assumption that is implicit in this principle. The assumption is that whenever the victims compare their experiences of pain at any two adjacent settings, they cannot distinguish any difference. So, for instance, if they first experience S_{20} and then S_{21} , they do not experience any change in the level of pain.¹⁰² Stating this assumption explicitly is important mostly to highlight that this is a feature of NTPCs that should be preserved in the final, weakened principle. Failing to preserve this assumption would lead to a rejection of NTPCs: if the victims can tell adjacent settings apart in terms of pain, then one could hardly claim that a single pressing of

¹⁰⁰ See also Barnett (2018: 6) and Nefsky (2011: 385; 2019: endnote 27) for similar claims.

¹⁰¹ For related discussions, see Spiekermann (2014: 79–80) and Regan (2000: 50–51).

¹⁰² I shall often talk about the victims comparing their experiences at different settings. Since only one setting can be directly experienced at a time, I shall assume that when the victims compare two settings, they experience one setting after the other in a reasonably short period of time.

the button does not make a difference to the victims' experiences. Let's add this newly unearthed assumption to the formulation of T:

Tolerance-2 (T2). *For all n , if the victims consider their experiences at settings S_n and S_{n+1} in comparison to each other, then if the victims experience no pain at S_n , then the victims experience no pain at S_{n+1} .*

Making this assumption explicit highlights an important idea but does not yet make much progress in addressing the paradox as no substantive modifications have been made so far. T2 at this stage collapses into T (Spiekermann 2014: 87).

Time for the second step. It is not enough that the victims cannot distinguish any two adjacent settings in terms of pain. They should also be able to distinguish their experiences at settings that are far enough apart. At the very least, they should be able to distinguish how they feel at S_0 from how they feel at S_{1000} . This claim, however, is not present in T or T2. This new claim is a substantive addition to the tolerance principle that weakens it. Let's formulate the principle with the new assumption:

Weak Tolerance (WT). *For all n , if the victims consider their experiences at settings S_n and S_{n+1} in comparison to each other, then if the victims experience no pain at S_n , then the victims experience no pain at S_{n+1} but they experience pain at S_{n+m} (where $m > 1$).*

Now, WT, unlike T, does not generate the paradox. If the victims experience no pain at S_0 and cannot tell the difference between their experience at any two adjacent settings, we are not forced to conclude that they also feel no pain at S_{1000} . It is still unclear, however, how it could be that while the victims experience no change in their level of pain when moving between any two adjacent settings, they do go from experiencing no pain at S_0 to experiencing excruciating pain at S_{1000} . One may even suggest that this second step is begging the question against those who accuse NTPCs

of generating the paradox. It seems arbitrary to assert without further argument that the victims' experiences change at some point.

Here is where the third step comes in. It consists in claiming that the “indistinguishable from” relation is nontransitive. For brevity, I shall the claim that the indistinguishability relation is nontransitive to the “Nontransitivity Claim” or NT.¹⁰³ It is then possible that the victims experience no change in the level of pain when moving between any two adjacent settings, but experience such a change when moving across several settings. Consider the following case as an illustration.

Pain Ex Nihilo. Suppose that at S_{10} the victims experience no pain. The button is pressed, and they move to S_{11} . By WT, they experience no change in the level of pain and cannot distinguish their experiences at S_{10} from their experiences at S_{11} . The button is pressed once again, and they move to S_{12} . They feel exactly as at S_{11} in terms of pain and cannot distinguish S_{11} from S_{12} in terms of pain. But now that they think about their previous experience, staying at S_{12} feels more painful than at S_{10} . So, they can distinguish the difference in pain between S_{10} and S_{12} .^{104, 105}

Here is a slightly more formal presentation of *Pain Ex Nihilo* assuming that “ \approx ” is “indistinguishable in terms of pain from”:

$$(A) (S_{10} \approx S_{11}) \wedge (S_{11} \approx S_{12}) \wedge \neg (S_{10} \approx S_{12})$$

¹⁰³ I shall further refer to it as “the indistinguishability relation” for convenience.

¹⁰⁴ The name of the case takes inspiration from Regan’s (2000: 50) term for the view that NTPCs are possible: “the ex nihilo claim”. The pain that the victims experience does not seem to increase at any single step, and so seems to appear ex nihilo — out of nowhere.

¹⁰⁵ If this seems counterintuitive, imagine that we are talking not about S_{10} and S_{12} but about something like S_1 and S_{900} . I am using the settings which are so close for clarity and brevity.

So, it is possible that while the victims cannot distinguish between any two adjacent settings in terms of pain, they can tell that their experiences change between two settings that are far enough apart. WT together with NT render cases like *Pain Ex Nihilo* possible. And it is plausible to treat these cases like instances of NTPCs. After all, the key feature of NTPCs is that no single act of the relevant type (e.g., pressing of the button) can make a morally significant difference. It seems plausible that the victims fail to distinguish any two adjacent settings in terms of pain precisely because a move from one setting to the adjacent one does not involve any change in the pain they experience. So, cases like *Pain Ex Nihilo* are possible even without accepting T. A weaker claim, WT, is sufficient (in conjunction with NT). And it is plausible to consider these cases to be NTPCs.

A further reason to believe that cases like *Pain Ex Nihilo* are possible is that they resemble phenomenal continua. A phenomenal continuum is a series of stimuli such that differences between any two adjacent stimuli are perceptually indistinguishable but where differences between stimuli which are far enough apart from one another in terms of differences in physical value are perceptually distinguishable (Raffman 2012: 313).¹⁰⁶ A common example of a phenomenal continuum is a series of coloured patches ranging from one shade of colour to an incompatible shade, for instance, from red to yellow. It is widely thought that phenomenal continua are possible. Moreover, many philosophers argue that claims like NT are necessary for accounting for the possibility of such continua.¹⁰⁷

Thus far, we have seen that the possibility of NTPCs can be defended if the T is modified into WT, and if we also accept NT.

¹⁰⁶ “Physical value” is the value of the physical parameter that is responsible for the changes in stimuli.

¹⁰⁷ See Fara (2001: 905–906) and Raffman (2012: 309–313) for overviews.

5.5 The Case Against NTPCs: Conceptual Work

In the remainder of the chapter, I aim to dismantle most of what I have presented in the previous section. Specifically, I shall argue that there is reason to think that NTPCs are impossible. And even if we grant that they are possible, they are unlikely to arise in practice.

In this section, I shall do three things. First, I offer an alternative explanation of how cases like *Pain Ex Nihilo* can arise, based on a version of what Diana Raffman (2012: 315) calls “the instability hypothesis”. This alternative explanation preserves the transitivity of the indistinguishability relation. Second, I argue that this transitivity does not force us to accept the paradoxical conclusion, contrary to what has been suggested above. Finally, I offer reasons to think why this alternative explanation provides a strong reason to think that NTPCs are impossible or at least significantly less prevalent than one may think.

5.5.1 *Transitivity Defended*

There is an alternative explanation of the possibility of cases like *Pain Ex Nihilo*. The key idea here is that an observer’s experiences of individual stimuli or objects in a phenomenal continuum change as she goes along the continuum.¹⁰⁸ Borrowing the term coined by Diana Raffman (2012: 315), I shall call this idea the “instability hypothesis” or IH. Let’s illustrate this idea with *Pain Ex Nihilo*. Recall that S_{10} is experienced as not painful. The next setting, S_{11} , if compared to S_{10} , is also experienced as not painful. But when S_{12} is compared to S_{11} , they are both experienced as slightly painful. And when S_{12} is compared to S_{10} , it is clear that S_{12} is experienced as more painful than S_{10} . Previously, the possibility of this case was explained by the claim that the indistinguishability relation is nontransitive. Now, in virtue of IH, an alternative explanation is available. Namely, S_{11} is experienced in one way when compared to S_{10} ,

¹⁰⁸ By “experience of stimuli” in this context, I mean “how stimuli look or feel” to an observer.

but when compared to S_{12} , it is experienced in a different way. So, if the victims choose to go from S_{12} back to S_{11} , they experience no change in pain (in accordance with WT) but experience the same slight pain that they felt upon reaching S_{12} . There is not one stable way, in which the victims experience S_{11} — their level of pain at S_{11} changes depending on which other settings they have experienced previously. As shall become clear from the discussion below, this alternative explanation based on IH, preserves the transitivity of the indistinguishability relation. I suggest calling the relevant case *Pain Ex Instability* (rather than *Pain Ex Nihilo*) when it is assumed that the change in the victims' experience is explained by IH rather than by NT.¹⁰⁹

Assuming IH, S_{11} gives rise to two distinct experiences. And then it makes sense to categorise it as “ S_{11} when compared to S_{10} ” (call this S_{11-}) and “ S_{11} when compared to S_{12} ” (call this S_{11+}). So, when S_{10} and S_{11} are considered in comparison to one another, they are indistinguishable in terms of pain and are both experienced as not painful. When S_{11} and S_{12} are considered in comparison to one another, they are also indistinguishable in terms of pain but are both experienced as slightly painful. S_{11-} is not painful and S_{11+} is slightly painful. If we accept IH, then a formalised version of *Pain Ex Instability* would then look like this:

$$(B) (S_{10} \approx S_{11-}) \wedge (S_{11+} \approx S_{12}) \wedge \neg (S_{10} \approx S_{12})$$

Contrast it with the formalisation of *Pain Ex Nihilo*:

$$(A) (S_{10} \approx S_{11}) \wedge (S_{11} \approx S_{12}) \wedge \neg (S_{10} \approx S_{12})$$

Transitivity would be violated if S_{11-} and S_{11+} were indistinguishable in terms of pain because then (B) would be identical to (A), where the relation is clearly nontransitive. However, S_{11-} and S_{11+} are not indistinguishable — in the present case they give rise to two genuinely distinct experiences. It cannot be inferred from (B) that

¹⁰⁹ I apologise to the reader for this abuse of the Latin language.

the indistinguishability relation is nontransitive. It would be similar to trying to infer that the relation “is the mother of” is nontransitive from the following example: Ann is the mother of Barbara. Beatrice is the mother of Claire. Ann is not the mother of Claire. Even though “is the mother of” is a nontransitive relation, we cannot make this inference from the example I have just mentioned, since Barbara and Beatrice are two different people (just like S_{11-} and S_{11+} are two distinct experiences). To claim that “is the mother of” is nontransitive, we need something like this: Ann is the mother of Beatrice. Beatrice is the mother of Claire. Ann is not the mother of Claire. If we could come up with an analogous example in the case of the indistinguishability relation, we would be able to infer that it is nontransitive. But assuming that IH is true, S_n does not have a stable way it is experienced in cases where S_{n-1} is experienced differently from S_{n+1} in terms of pain. In other words, assuming IH is true, settings like S_{11} would always give rise to two distinct experiences such as S_{11-} and S_{11+} . Thus, there would be no cases from which we would be able to infer that the indistinguishability relation is nontransitive.¹¹⁰ Any case involving perceptual indistinguishability of very similar stimuli would be analogous to (B) rather than to (A).

A defence of transitivity similar to this one is carried out by John Broome (1991: 100–103). Although Broome argues for the transitivity axiom for rational preferences and not for the indistinguishability relation, his argument can be used as a nice analogy to my discussion here.

¹¹⁰ Generally, we only need one instance of nontransitivity of a relation to infer that the relation is nontransitive. No matter how many cases involving Barbaras and Beatrices we can come up with, just one counterexample involving only Ann, Beatrice, and Claire is sufficient for demonstrating nontransitivity. This could work similarly for the indistinguishability relation. Perhaps there is an example of its nontransitivity which I haven’t discussed? I doubt it, but I do not need to insist that such an example cannot be found. All that is needed for my present purposes is that we can never infer that the relation is nontransitive in the context of phenomenal continua (even if we can so infer in some other context). And if the instability hypothesis is true, then we indeed cannot make that inference. If the relation is nontransitive in some other context, then when I say that it is transitive, this would not be, strictly speaking, true — but importantly, this would not affect the cogency of my argument.

Broome (1991: 101–102) designs the following scenario. A person, call him Maurice, is in a decision situation with three options: to go to Rome (R), to go mountaineering in the Alps (M), to stay at home (H). He prefers R to M because visiting Rome is more culturally enriching than climbing the mountains. He prefers H to R because sightseeing is boring for him. He prefers M to H because staying at home rather than going to the mountains is cowardly. There is thus an apparent intransitivity in Maurice’s preferences. However, Broome argues, this intransitivity can be dealt with. The transitivity of Maurice’s preferences is restored if H is separated into two different outcomes: H_R , which is staying at home without having turned down the mountaineering trip; and H_M , which is staying at home having turned down the mountaineering trip. H_M is cowardly and H_R is not. So, Maurice prefers R to M, and H_R to R. But if he has a choice between M and H, there is H_M instead of H. And he prefers M to H_M without violating transitivity. What about H_R and M? It is impossible that Maurice encounters a choice between these two alternatives because H_R is an alternative of staying at home without rejecting M. And if he chooses H_R , he rejects M. So, the choice situation between H_R and M is impossible. In sum, the transitivity of Maurice’s preferences is preserved by separating H into H_R and H_M .

Just as H is divided into H_R and H_M in Broome’s argument, if we assume that the instability hypothesis holds, S_{11} is divided into S_{11-} and S_{11+} in *Pain Ex Instability*. In Broome’s argument, this division is what makes it possible to preserve the transitivity of preferences and in *Pain Ex Instability* — the transitivity of the “indistinguishable from in terms of pain” relation. However, there are also several significant differences between these two cases. A discussion of these differences will help to grasp how the instability hypothesis works and why it is plausible.

The first difference is that H_R and H_M are different outcomes between which an agent can have a preference. Maurice, for instance, prefers H_R to H_M . S_{11-} and S_{11+} , on the contrary, are not separate settings on the machine. They are just two ways in which one setting could be experienced. What is the relation between S_{11-} and S_{11+} ? Since the former is experienced as not painful and the latter is experienced as slightly painful, they are distinguishable in terms of pain. This is not surprising — the relation between

S_{11-} and S_{11+} is the same as the relation between S_{10} and S_{12} . This fact does not undermine the instability hypothesis. The fact that S_{11+} is more painful than S_{11-} does not mean that the victims can compare how S_{11} feels in relation to itself and find that it is more or less painful than itself. Whenever they assess how much pain they feel at S_{11} , they will experience either S_{11-} (no pain) or S_{11+} (slight pain) depending on which setting they are comparing S_{11} to.

The second difference is related to a particular worry Broome has about fine-grained individuation of outcomes. He claims that if there were no limits to fine-graining, then all preferences can be shown to be transitive — as long as the outcomes are individuated in a sufficiently fine-grained way (Broome 1991: 102). Letting transitivity to be a vacuous requirement on preferences is unacceptable, and Broome suggests two ways to avoid it (he prefers the second way). The first one is to prohibit fine-graining when an agent cannot have a preference between the two potentially different outcomes. The second way is to impose no limits on fine-graining but to formulate a rationality requirement of indifference for preferences between certain outcomes. So, agents rationally ought to be indifferent between certain outcomes. If Maurice rationally ought to be indifferent between H_R and H_M , his preferences would not be transitive. By introducing an indifference constraint, Broome attempts to save the requirement of transitivity from being empty.

Does this worry about the emptiness of the transitivity requirement translate to the case of the instability hypothesis? It does not. An observer cannot claim that she experiences a certain stimulus in more than one distinct way just because she wishes to. So, S_{11} cannot be divided into S_{11-} and S_{11+} unless they constitute genuinely distinct ways S_{11} is experienced. And whether to experience a certain stimulus in a certain way is not up to the observer. This is unlike the case with preferences between outcomes, where it is up to the agent to individuate the outcomes as finely as she wants and then to form preferences between the individuated outcomes (unless she accepts the indifference constraint). The point is that there is no need for any analogue of the indifference constraint in the case of the instability hypothesis because there is no danger of the transitivity of the indistinguishability relation becoming vacuous. There

is no danger because, once again, the limit to individuating the way a stimulus is experienced is provided naturally — it is an empirical question whether a certain stimulus gives rise to more than one genuinely distinct experience.¹¹¹

5.5.2 *Instability and Transitivity Without Paradox*

I have argued that the indistinguishability relation is transitive if IH is accepted. I shall now explain why this does not lead to the paradoxical conclusion and why this is consistent with the possibility of phenomenal continua. Recall why one may think that transitivity is problematic. It is because if a setting S_{n-1} is perceptually indistinguishable from S_n , and S_n from S_{n+1} , then we seem to be forced to conclude that S_{n-1} is indistinguishable from S_{n+1} . But once we accept IH, this worry is no longer relevant. It is irrelevant because the victims can distinguish two distinct experiences, S_n and S_{n+} , to which S_n gives rise. So, when they compare S_{n-1} with S_n , they are actually comparing S_{n-1} with S_n , and conversely, when they compare S_{n+1} with S_n , they are actually comparing S_{n+1} with S_{n+} . We are not forced to conclude that S_{n-1} is indistinguishable from S_{n+1} . This implies that victims can experience non-adjacent settings as distinguishable in terms of pain. And that is all we need to avoid the paradoxical conclusion.

Similarly, we can see that preserving transitivity by accepting IH does not rule out the possibility of phenomenal continua. It is possible that the victims would be unable to perceive the difference in pain between any two adjacent settings. This is because when they move from one setting to an adjacent one, the way they experience the latter changes (as was the case with S_{11} in *Pain Ex Instability*). But, as I have just shown in the previous paragraph, the victims can distinguish between settings that are far enough apart from each other. So, we have both the key features of a phenomenal continuum.

All that I have said in the above two paragraphs can also be expressed in a much briefer, but less explicit way. WT, which my argument accepts, is already sufficient for

¹¹¹ In section 5.6, I shall offer some empirical evidence in support of the instability hypothesis.

both avoiding the paradox and ensuring the possibility of phenomenal continua. Here is a reminder of how WT is formulated:

Weak Tolerance (WT). For all n , if the victims consider their experiences at settings S_n and S_{n+1} in comparison to each other, then if the victims experience no pain at S_n , then the victims experience no pain at S_{n+1} but they experience pain at S_{n+m} (where $m > 1$).

A defender of the NTPCs could argue that WT can only be made sense of if the indistinguishability relation is nontransitive. But IH makes it possible to accept both WT and the claim that the relation is transitive. That way, we avoid the paradox, and preserve the possibility of both phenomenal continua and the transitivity of the relation.

5.5.3 *Instability and Transitivity Against the Possibility of NTPCs*

Let's turn to the question of how accepting WT and IH helps us see that NTPCs are impossible (or at the very least quite rare). The key point I advance in this subsection is that if we assume that collective impact cases such as the *Harmless Torturers* are like *Pain Ex Instability* rather than *Pain Ex Nihilo*, then there is a good reason to think that an individual pressing of the button would at least have a chance of increasing the victims' levels of pain. And if it is the case that individual pressings of the button could make a difference, then the cases in question are most often *not* NTPCs. That is because a distinctive feature of an NTPC is that a single act of the relevant type *never* makes a morally significant difference. Put differently, cases like *Pain Ex Instability* would be threshold cases.

So, why would cases like *Pain Ex Instability* be threshold cases? The answer is that in such cases, at least some increases in pain would be traceable to individual acts of pressing the button. The key obstacle to tracing is the fact that victims cannot distinguish any two adjacent settings in terms of pain. If so, how could we claim that at any given step a single pressing of the button causes an increase in the level of pain?

This problem leads us to conclude that in *Pain Ex Nihilo* no single pressing makes a difference. However, we need not make this conclusion in *Pain Ex Instability*. In this latter case, due to the IH, there is a plausible explanation of how the level of pain can increase due to a single pressing despite the victims not being able to tell the difference between any two adjacent settings.

Consider *Pain Ex Instability* again. Once the victims reach setting S_{11} after S_{10} , they experience no pain and feel no difference between S_{10} and S_{11} . Once they reach S_{12} , they judge that S_{12} is more painful than S_{10} , yet S_{11} is experienced no different from S_{12} . The explanation is that the victims are unable to experience the difference between S_{11} and S_{12} because their experience of S_{11} shifted from S_{11-} to S_{11+} once they arrived at S_{12} . They do not notice the shift itself. Moreover, if they go back to S_{11} after S_{12} , S_{11} would be experienced as slightly painful, as if they have not undergone any shift in experience. However, since we know how the IH works, we may argue that *if the victims had not moved from S_{11} to S_{12} , they would not have experienced the slight pain that they experience at S_{12}* . The fact that S_{11} and S_{12} are indistinguishable does not entail that S_{11} was painful prior to the move from S_{11} to S_{12} . This is precisely because the reason why S_{11} is experienced as slightly painful is that a move has been made to S_{12} , which shifted the way S_{11} is experienced. So, it is plausible that S_{11} was experienced as not painful prior to moving to S_{12} , but once the button is pressed, the setting switches to S_{12} , the level of pain increases, and the way S_{11} is experienced shifts. The pressing of the button could plausibly be the difference-maker or at least among the difference-makers both to the increase in pain and to the shift in how S_{11} is experienced.

If we assume that IH is true, could there be an alternative explanation of how the overall level of pain could increase while no two adjacent settings are distinguishable? If there were an alternative explanation that assumed the truth of IH but did not make tracing the change in the level of pain to a single pressing of the button possible, then the case against NTPCs would be significantly weaker. Some may propose that had the button not been pressed at S_{11} , the change in the level of pain would still have occurred, so S_{11} could already be experienced as slightly painful. But how could S_{11} be

experienced as slightly painful if no setting ahead, such as S_{12} , has been tried? The change in the level of pain could occur, for instance, due to some change in the environment or due to some feature of how the victims' brains work, but not directly due to the pressing of the button (Nefsky 2011: 390–392). It must be then, that when S_{11} is considered in comparison with S_{10} , it is indistinguishable from this setting and is experienced as not painful. But then, without any further intervention, a shift in how S_{11} is experienced occurs, and it starts to be experienced as slightly painful. If, after this shift, the victims would go back to S_{10} , S_{10} might be experienced as slightly painful as well.

This alternative explanation does not strike me as plausible. It seems more natural to say that a change in how a certain setting is experienced is caused by specific moves along the continuum. After all, if a certain setting, as in the previous example, S_{11} is already considered, and is experienced in a certain way, why would it undergo a shift in how it is experienced if no further intervention is made? Moreover, how could such a shift undergo in a way that preserves the continuous progression characteristic of phenomenal series such as the one in *Harmless Torturers*? It seems to me more plausible to claim that if S_{11} is experienced as slightly painful even though S_{10} was experienced as not painful, then the shift occurs in how S_{10} is experienced *because* the next setting, S_{11} , is considered.

I struggle to think of other alternative explanations. Note that any other possible explanation that might rely on the claim that the indistinguishability is nontransitive is not acceptable, since we are proceeding on the assumption that IH holds and that the relation is thus transitive.

Thus, assuming IH, we have a good reason to believe that the relevant cases are like *Pain Ex Instability* rather than *Pain Ex Nihilo*, and that whenever there is an increase in the level of pain, it is potentially traceable to a certain pressing of the button.

5.5.4 Instability and Transitivity Against the Prevalence of NTPCs

However, suppose that for the sake of the argument we grant that the alternative explanation that the effective difference-making causes of some increases in pain may lie in some changes in the environment or in some neural processes. If we grant this, my argument does not rule out the possibility that *some* increases in pain are not caused by any single pressing of the button (even if IH is true). It could be, for instance, that in two moves (i) S_{20} to S_{21} and (ii) S_{21} to S_{22} the level of pain increases. However, had the button not been pressed at S_{21} , the change in the level of pain would still have occurred. This would establish that NTPCs are possible, but not that they actually occur.

My argument, however, provides a reason to think that in virtue of the availability of an IH-based explanation, *most* cases would not be NTPCs (even if some of them *could* be NTPCs). After all, note that I do not need to claim that in *Harmless Torturers*, each pressing of the button causes an increase in the level of pain. I do not even need to claim that some of the pressings do that. To claim that a case is not an NTPC, it is sufficient that *at least one pressing of the button has at least a chance of making a difference* to the pain experienced by the victims. And it seems to me overwhelmingly plausible that this weak condition would be met in most cases. Indeed, it seems peculiar to deny that a single pressing of the button has not even a chance of making a difference to the pain that the victims experience. What makes it especially peculiar is that the fact that the victims do not experience the difference in pain between two adjacent settings does not lend support to the claim that the level of pain has not increased in the step between these two settings. Assuming IH, it may very well increase, and quite plausibly due to a certain pressing of the button.

A point made by Nefsky (2011: 392–394) can be considered as an objection to my reasoning. The point is that even if we accept that in cases like *Harmless Torturers* at least one pressing of the button has at least a chance of making a difference, these cases are still not threshold or triggering cases, but rather quasi-threshold or quasi-triggering ones. A quasi-triggering case is different from a genuinely triggering one in that while some acts like pressing the button *could* make a difference, it does not follow

that an act of this kind *does* make a difference. In a genuinely triggering case, by contrast, at least one act of the relevant kind *does* in fact make a morally significant difference. So, while at least one pressing of the button could make a difference, it could also turn out that each increase in pain is in fact caused by some factor unrelated to pressing the button.

My response to this is that a quasi-threshold case is still not a non-threshold case. And the difference between these two kinds of cases is already significant. So, a conclusion that most cases are of the threshold or quasi-threshold variation would be significant. The key reason is that the inefficacy problem is much more troublesome if we face a non-threshold case rather than a quasi-threshold case. For example, in quasi-threshold cases, it is not too problematic to classify pressings of the button at least as partial causes of the increases in pain provided that there is at least a chance of their making a difference. However, in the case of non-threshold cases, doing so would be very problematic, at least according to standard difference-making views of causation. After all, it is the feature of non-threshold cases that no act of the relevant kinds makes a difference. There are other examples, but even the one I've mentioned would suffice.¹¹² Showing that cases like *Harmless Torturers* are threshold or at least quasi-threshold ones is significant.

I have argued that even if we admit the possibility of NTPCs (which we are not rationally required to do), it is reasonable to think that they would be much less prevalent than threshold or at least quasi-threshold cases.

5.5.5 *Taking Stock*

I have argued that accepting the instability hypothesis (IH) helps us defend the transitivity of the indistinguishability relation. This, in turn, enables us to argue that collective harm cases like the *Harmless Torturers* are in fact akin to the *Pain Ex Instability* case, rather than the *Pain Ex Nihilo*. And in cases like *Pain Ex Instability*,

¹¹² See Tiefensee (2022) for a discussion of different problems that non-threshold cases raise and that are not relevant for threshold cases.

we are able to explain of how the level of pain increases by tracing the increase to specific pressings of the button. This implies that cases like *Pain Ex Instability* are threshold cases. However, if we grant that *Pain Ex Instability* admit of an alternative, non-tracing explanations, then such cases could be non-threshold ones. But even this is granted, NTPCs would be much less prevalent than threshold or non-threshold cases. Importantly, the discussion in this section has proceeded on the assumption that IH holds. The next section, drawing on Raffman's (2012) work, provides empirical backing for this hypothesis.

5.6 The Case Against NTPCs: Empirical Work

Raffman (2012: 315–318) has carried out an experiment that provides empirical support for the instability hypothesis.¹¹³ The experiment is as follows.¹¹⁴

There is a phenomenal continuum of 41 coloured patches from one shade of green to a noticeably different shade of green. Let P_n be a coloured patch, where n is the number of the patch in the series from 1 to 41. Some adjacent patches are identical, and some are slightly different in hue, but below the discrimination threshold of any of the test subjects. At all times during the experiment, the subjects can see all 41 patches that are arranged in a circle. Whenever subject is asked to look at any two specific patches, they are highlighted by black dots appearing near them. This stage of the experiment can be seen in figure 1 below. The experiment begins with a comparison of two adjacent patches, say P_1 and P_2 . If the subject judges them to look different, then the next pair is offered, in this case, P_2 and P_3 . If the subject judges a pair to look the same, then she is offered a different patch, which is not a patch from the series, call it the D patch. The D patch can be seen inside the circle in figure 2

¹¹³ Raffman has designed and run the experiment in collaboration with Delwin Lindsey and Angela Brown.

¹¹⁴ My presentation of the experiment differs slightly from the presentation made by Raffman (2012). All the significant details are kept the same.

below. The starting colour of the D patch is randomised. It is still green, but of a random hue, so it most often appears clearly distinct from any of the P patches. The subject adjusts the colour of the D patch until it matches the colour of the pair under consideration. The subject has to decide for herself when she is happy with the adjustment. The adjustment is carried out on a computer by moving the computer mouse. After the subject has adjusted the colour of the D patch to the colour of the pair under consideration, the D patch disappears, and the next pair (P_2 and P_3) is considered. This goes on until the subject reaches the final pair in the series. Here is an example of how a part of the experiment could go:

Step 1: The subject compares P_1 and P_2 . The subject judges them “indistinguishable”.

Step 2: The subject observes the D patch and adjusts its colour until it matches the colour of P_1 and P_2 .

Step 3: The subject compares P_2 and P_3 . The subject judges them “indistinguishable”.

Step 4: The subject observes the D patch and adjusts its colour until it matches the colour of P_2 and P_3 .

Figure 1





Figure 2



Both figure 1 and figure 2 are taken from Raffman (2012: 316)

The results of the experiment show that for a vast majority of pairwise comparisons between the P patches, the subjects judged the patches to look the same (Raffman 2012: 316–317). However, the hue assigned to the D patches by the subjects progressed systematically with the differences in physical values of the patches. So, for instance, the subjects assigned a different colour hue to the D patch when adjusting it to P_1/P_2 , than they did when adjusting it to the pair P_2/P_3 . So, the patch P_2 was matched to a different hue when compared to P_1 than when it was compared to P_3 . To put the example differently, suppose that P_1 is indistinguishable from P_2 , and the D patch matched to them has the value x . When P_2 is compared to P_3 , they are also indistinguishable in terms of colour. However, the D patch, when adjusted to match the colour of P_2 and P_3 , has the physical value $x+1$. In sum, while the subjects were unable to detect any differences between adjacent patches, they assigned a different hue to the D patch when matching it with different pairs.

The results can be interpreted to suggest that one and the same patch is experienced differently by the subjects depending on which patch it is compared to. When P_2 is compared to P_1 , they are perceptually indistinguishable. So are P_2 and P_3 . But the subjects assign a different hue to the D patch when matching it to the P_1/P_2 pair than when matching it to the P_2/P_3 pair. The explanation for assigning different hues to the D patch can be that P_2 is experienced differently in comparison to P_1 than in comparison to P_3 . This way, the instability hypothesis is shown to have empirical support. Admittedly, this interpretation of the results is not conclusive. The subjects, for instance, did not experience that the hue shifted while they were going through the series. However, the progression in the hue of the D patch, together with the fact that the subjects could not distinguish between any two P patches in a pairwise comparison do suggest that an explanation of these results in terms of the instability hypothesis is plausible.

The point of this section has been to present the results of an experiment that lends support to the instability hypothesis. Most of the argument in the previous section is built on the assumption that IH holds. We have seen that the experimental results can be interpreted as supporting the hypothesis.

One may question the relevance of these results for the argument in this chapter, given that the experiment is run in the context of colour perception, while the chapter primarily focuses on pain. I do not claim that pain and colour perception work the same. But I do not need to claim that. All I needed from the discussion of the experiment is the instability hypothesis. And this hypothesis is most relevant for phenomenal continua. Indeed, it describes a feature of observers' perception in these continua. Phenomenal continua seem clearly possible in the case of colour perception. But it seems also plausible that continua of a similar kind are possible in the case of pain perception. *Harmless Torturers* contains such a continuum, and it does not seem to be a practically impossible case. And since phenomenal continua appear possible both in the case of phenomenal experiences of colour and in the case of phenomenal experiences of pain, there is reason to think that the instability hypothesis would hold in both cases.

5.7 Conclusion

In this chapter, I have argued that there is a good reason to think that purely phenomenal non-threshold collective impact cases cannot possibly arise (or at least that they are likely to be very rare). I have started by making the observation that non-threshold cases seem to generate the sorites paradox. However, there are arguments showing that they do not generate it. One such argument that applies specifically to NTPCs involves accepting that the indistinguishability relation is nontransitive. If this claim is accepted, then we could argue that even though no two stimuli very close in terms of the relevant physical value could be perceptually distinguished, stimuli that are sufficiently far distinct in terms of physical value, can be distinguished. Thus, the paradox is avoided but the possibility of NTPCs is preserved: it could still be the case that no single step between the stimuli causes any change in the relevant experiences. But, drawing on Raffman's work (2012), I have suggested that we should accept an alternative strategy of addressing the paradox. This alternative strategy relies on what has been called the 'instability hypothesis' and it preserves the transitivity of the

relation. This strategy also has empirical support — it seems accurate in describing how our perception in phenomenal continua works. Moreover, assuming that the instability hypothesis holds also helps us argue that NTPCs are impossible, or at least rare. This result is significant. If these cases are impossible, then our lack of solution to the problems they give rise to, such as particularly troubling versions of the inefficacy problem, is not that worrying.

This is not to say that the inefficacy problem completely ceases to be worrying. After all, even threshold cases face versions of this problem (Budolfson 2018, Nefsky 2019, Tiefensee 2022). And there are not purely phenomenal non-threshold cases too. They generate versions of the problem as well, and my argument says nothing against the existence of such cases.

The reason why my argument does not apply to these cases is that the instability hypothesis is a claim about our psychology. If my argument is correct, it helps us explain why we can often trace changes to our phenomenal experiences of pain to specific interventions. But in many non-threshold cases whether we can perceive the differences in the morally relevant aspects is irrelevant. Recall *Fishermen* from the previous chapter. In that case, it seems irrelevant whether we can perceive or observe the changes caused by a single instance of using a big net to the fish population's ability to replenish. What is important for the question of whether *Fishermen* is a non-threshold case is whether the ability of fish population to replenish is sensitive to changes caused by an instance of using a big net. The argument I put forward in this chapter does not say anything about the possibility of insensitivities of this sort. The possibility of non-threshold cases of a not purely phenomenal kind is a problem for future work.

Chapter 6

Concluding Remarks

As the concept of imperfect duty gains prominence and attracts increased attention from philosophers, it becomes more widely applied throughout various debates, often with promising results. However, with prominence and wider application grow misconceptions and confusions. Furthermore, underdiscussed assumptions become more ingrained as more theorists make them without questioning. My main aim in this dissertation has been to intervene in three of the debates surrounding the concept of imperfect duty to dissolve some of the confusions, clarify some of the misconceptions, and unearth and evaluate some of the underdiscussed assumptions. In this concluding chapter, I summarise the main results achieved in this dissertation and outline a direction for further research.

Chapter 2 addressed the debate about the relationship between Kantian imperfect duties and supererogation. It has been widely accepted that a moral theory that includes Kantian imperfect duties has no space for the supererogatory at least in part due to their limitlessness. This view, which I call the ‘Incompatibility Thesis’, should be rejected. There are two reasons why the grounds for rejecting the Incompatibility Thesis have not been sufficiently clear. First, the proponents of the Incompatibility Thesis have generally failed to recognise the importance of the distinction between weak and strong forms of supererogation for this debate. Second, they have not distinguished the different features constitutive of limitlessness. Once we make the relevant distinctions, the reasons for rejecting the Incompatibility Thesis become much clearer.

I argued that on any of the three plausible interpretations of limitlessness (persistence, open-endedness, and requiredness), limitless imperfect duties are compatible at least with weak supererogation. This is a significant result. First, only limitlessness as requiredness is incompatible with strong supererogation. Second,

weak supererogation is sufficient to explain the relevant moral phenomena, such as heroic actions that are very good to perform but are not strictly required. Lack of compatibility with strong supererogation is not a significant drawback. My discussion in Chapter 2 thus makes a twofold contribution to the literature. It enhances our understanding of the nature of limitlessness of imperfect duties. It also shows that the Incompatibility Thesis should be rejected. We can have a moral theory with end-based imperfect duties and supererogation. And even if Kant's ethics ultimately would turn out to have no conceptual space for supererogation, it is not because of the nature of imperfect duties.

Chapter 3 questioned the widely held assumption that there is an interesting connection between latitude permitted by imperfect duties and demandingness of morality. It has been assumed that latitude permitted by imperfect duties 'limits' the demandingness of morality. The aim of this chapter was to argue that there is no interesting sense in which this is true. The first task was to identify the plausible interpretations of this connection as they have not been articulated sufficiently clearly in the literature. I have identified three alternative interpretations. The presence of imperfect duties that permit extensive latitude (i) *explains* why morality is moderately demanding, (ii) *implies* that morality is moderately demanding *without explaining* why this is so, or (iii) *is merely compatible* with moderately demanding morality. The second task was to determine which of the interpretations are acceptable. I have argued that the first two interpretations should be rejected or set aside as theoretically uninteresting because they either involve false claims, make question-begging assumptions, or state trivial truths. The third interpretation is the one we should accept. However, the connection between latitude and demandingness on this interpretation is not very interesting. All that we can conclude from it is that a moderately demanding morality that includes imperfect duties with extensive latitude *is possible*. But whether morality is moderately demanding depends on further arguments, independent of claims about latitude permitted by imperfect duties.

According to some recent discussions, imperfect duties are promising as bases for a solution to the inefficacy problem that arises in collective harm cases (Albertzart

2019; Baatz 2014; Nefsky 2021). It is thought that contributions of individual agents to collectively caused harmful outcomes could be condemned precisely because these contributions violate imperfect duties (even if they do not make a difference to these outcomes). In Chapter 4, I argued that the Kantian end-based account of imperfect duties is not helpful in this respect. The key problem is that it is ultimately very unlikely that an agent's pattern of behaviour in any given collective harm case would be the reason why this agent violates her imperfect duty to adopt the end of others' happiness. This entails that we would very rarely be able to condemn individual agents' actions in collective harm cases on the grounds of their violating imperfect duties, even when condemning such actions would be intuitively correct.

This result is not entirely negative. There is a reason to think that despite not being useful for addressing the *theoretical* puzzle of the inefficacy problem, imperfect duties can be helpful in addressing the *practical* side of the problem. Specifically, it is a plausible conjecture that they can guide agents to communicate, cooperate, and put in place institutional measures that will help them avoid collectively causing harmful outcomes.

Finally, Chapter 5 took a detour from the main object of discussion to consider collective impact cases in more detail. Its aim was to argue that there is a good reason to think that at least a certain kind of cases that generate particularly problematic versions of the inefficacy problem are either impossible or very unlikely to arise. Specifically, it discussed purely phenomenal cases, that is, cases where harm in question is purely experiential. I have argued that there is an explanation that shows that a single act of the relevant kind has at least a chance of making a difference to experienced harm and that it may be possible to trace increases in harm to acts of this kind. This explanation rests on a certain empirical assumption about how our ability to discriminate between perceptual stimuli works. I have drawn on Raffman's (2012) empirical research to support this assumption. If my argument in this chapter is sound, then there is a good reason to think that non-threshold cases at least of purely phenomenal variety are impossible (or that such cases arise very rarely). This result

does not completely remove the challenges posed by the inefficacy problem, but it makes the problem at least somewhat less worrying.

The results achieved in this dissertation are largely negative in the sense that they are to a significant extent directed against ideas or views held in the debates surrounding imperfect duties. However, these results also open a promising direction for further research.

When presenting the Kantian end-based account in the Introduction, I set several important substantive questions aside. Specifically, I have not addressed the questions of what makes an end obligatory, under what exact conditions we can say that we succeed in adopting and promoting an end, and how we should weigh our obligatory ends against our discretionary ends. And if we take some end-based view on imperfect duties seriously and aim to develop a satisfactory comprehensive account of such a view, these are the questions we shall need to address. My dissertation can be seen as doing important conceptual work that serves as a ground for this further investigation of substantive issues. Addressing these issues will help us make progress in some of the debates that I have engaged with in this dissertation.

For instance, in Chapter 3, I argued that the conceptual connection between latitude permitted by imperfect duties and demandingness of morality is of little theoretical interest. But this negative result also suggests that whether morality that includes imperfect duties is moderately demanding is to a significant extent a question that depends on our substantive views about the relevant obligatory ends. By establishing what conditions are relevant for determining when we succeed in adopting an obligatory end and what counts as promoting it to a sufficient extent, we would have a clearer picture of the costs involved. We could then make progress on the question of how demanding would morality that includes imperfect duties be.

In Chapter 4 we have seen that Kantian imperfect duties do not help in addressing the inefficacy problem. However, I made a tentative suggestion that they may still help address the practical side of the problem: they *may* guide us to avoid the harmful outcomes in collective impact cases. Whether they *do* offer such guidance remains to be seen. This is another example of further research on substantive aspects of

imperfect duties that can be carried out using the results that have been achieved in this dissertation.

These are just a few examples of debates where a fully developed end-based account of imperfect duties could fruitfully build on the work done in this dissertation. I do not suggest that all the important conceptual work that could underlie answers to the further substantive questions is done. To be sure, there are more misconceptions to correct and more problematic assumptions to bring to light. Nevertheless, this dissertation made some progress on this task.

Bibliography

- Albertzart, Maïke. 2019. "A Kantian Solution to the Problem of Imperceptible Differences." *European Journal of Philosophy*. 27 (4): 837–851.
- Archer, Alfred. 2016. "Are Acts of Supererogation Always Praiseworthy?" *Theoria*. 82 (3): 238–255.
- Archer, Alfred. 2018. "Supererogation." *Philosophy Compass*. 13 (3). doi:10.1111/phc3.12476.
- Arntzenius, Frank and McCarthy, David. 1997. "Self-Torture and Group Beneficence." *Erkenntnis* 47 (1): 129–144.
- Atterton, Peter. 2007. "A Duty to Be Charitable? A Rigoristic Reading of Kant." *Kant-Studien*. 98 (2): 135–155.
- Baatz, Christian. 2014. "Climate Change and Individual Duties to Reduce GHG Emissions." *Ethics, Policy and Environment*. 17 (1): 1–19.
- Baatz, Christian and Voget-Kleschin, Lieske. 2019. "Individual Contributions to Harmful Climate Change: The Fair Share Argument Restated." *Journal of Agricultural and Environmental Ethics*. 32 (4): 569–590.
- Barnett, Zach. 2018. "No Free Lunch: The Significance of Tiny Contributions." *Analysis*. 78 (1): 3–13.
- Baron, Marcia. 1995. *Kantian Ethics Almost Without Apology*. Ithaca, New York: Cornell University Press.
- Baron, Marcia. 2015. "The Supererogatory and Kant's Imperfect Duties." In M. Timmons and R.N. Johnson (eds.) *Reason, Value, and Respect: Kantian Themes from the Philosophy of Thomas E. Hill, Jr.* pp. 215–231. Oxford: Oxford University Press.
- Baron, Marcia. 2016. "A Kantian Take on the Supererogatory." *Journal of Applied Philosophy*. 33 (4): 347–362.
- Berg, Amy. Forthcoming. "Is There a Duty to Read the News?" *Journal of Moral Philosophy*.

- Biss, Mavis. 2015. "Kantian Moral Striving." *Kantian Review*. 20 (1): 1–23.
- Bratman, Michael. 2014. *Shared Agency: A Planning Theory of Acting Together*. New York: Oxford University Press.
- Brock, Dan. 1991. "Defending Moral Options." *Philosophy and Phenomenological Research*. 51 (4): 909–913.
- Broome, John. 1991. *Weighing Goods: Equality, Uncertainty and Time*. Oxford: Basil Blackwell.
- Budolfson, Mark Bryant. 2018. "The Inefficacy Objection to Consequentialism and The Problem with The Expected Consequence Response." *Philosophical Studies*. 176 (7): 1711–1724.
- Budolfson, Mark Bryant. (unpublished manuscript). "The Inefficacy Objection to Deontology: What It Is, Why It Is Important, and How to Respond to It."
- Chappell, Richard Yetter. (unpublished manuscript). "There is No Problem of Collective Harm."
- Cohen, G.A. 2000. *If You're an Egalitarian, How Come You're So Rich?* Cambridge, MA: Harvard University Press.
- Crisp, Roger. 2013. "Supererogation and Virtue." In Timmons, Mark (ed.), *Oxford Studies in Normative Ethics, Volume 3*, pp. 13–34. Oxford: Oxford University Press.
- Cummiskey, David. 1996. *Kantian Consequentialism*. New York: Oxford University Press.
- Dancy, Jonathan. 1993. *Moral Reasons*. Oxford: Blackwell.
- Denis, Lara. 2001. *Moral Self-Regard: Duties to Oneself in Kant's Moral Theory*. London: Routledge.
- Dorsey, Dale. 2013. "The Supererogatory, and How to Accommodate It." *Utilitas*. 25(3): 355–382.
- Dougherty, Tom. 2016. "Moral Indeterminacy, Normative Powers and Convention." *Ratio*. 29 (4): 448–465.
- Dougherty, Tom. 2017. "Altruism and Ambition in the Dynamic Moral Life." *Australasian Journal of Philosophy*. 95 (4): 716–729.

- Elson, Luke. 2016. "Tenenbaum and Raffman on Vague Projects, the Self-Torturer, and the Sorites." *Ethics*. 126 (2): 474–488.
- Fanciullo, James. 2021. "The Psychological Basis of Collective Action." *Philosophical Studies*. 178 (2): 427–444.
- Fara, Delia Graff. 2001. "Phenomenal Continua and the Sorites." *Mind*. 110 (440): 905–935.
- Feinberg, Joel. 1960. "Supererogation and Rules." *Ethics*. 71 (4): 276–288.
- Flescher, Andrew M. 2003. *Heroes, Saints, and Ordinary Morality*. Washington, D.C.: Georgetown University Press.
- Gert, Joshua. 2016. "The Distinction Between Justifying and Requiring: Nothing to Fear." In E. Lord and B. Maguire (eds.) *Weighing Reasons*, pp. 157–172. Oxford: Oxford University Press.
- Greenspan, Patricia. 2010. "Making Room for Options: Moral Reasons, Imperfect Duties, and Choice." *Social Philosophy and Policy*. 27 (2): 181–205.
- Gregor, Mary J. 1963. *Laws of Freedom: A Study of Kant's Method of Applying the Categorical Imperative in the "Metaphysik der Sitten"*. Oxford: Blackwell.
- Guevara, Daniel. 1999. "The Impossibility of Supererogation in Kant's Moral Theory." *Philosophy and Phenomenological Research*. 59 (3): 593–624.
- Hale, Susan. 1991. "Against Supererogation." *American Philosophical Quarterly*. 28 (4): 273–285.
- Hanser, Matthew. 2014. "Imperfect aiding." In S. Luper (ed.), *The Cambridge Companion to Life and Death*, pp. 300–315. New York: Cambridge University Press.
- Herman, Barbara. 2007. *Moral Literacy*. Cambridge, Massachusetts: Harvard University Press.
- Herman, Barbara. 2012. "Being Helped and Being Grateful: Imperfect Duties, the Ethics of Possession, and the Unity of Morality." *Journal of Philosophy* 109 (5/6): 391–411.
- Herman, Barbara. 2021. *The Moral Habitat*. Oxford: Oxford University Press.

- Heyd, David. "Supererogation." In E. N. Zalta (ed.) *The Stanford Encyclopedia of Philosophy* (Winter 2019 Edition), <<https://plato.stanford.edu/archives/win2019/entries/supererogation/>>.
- Hill, Thomas E. 1971. "Kant on Imperfect Duty and Supererogation." *Kant-Studien*. 62: 55–76.
- Hill, Thomas E. 2002. *Human Welfare and Moral Worth*. New York: Oxford University Press.
- Hiller, Avram. 2011. "Climate Change and Individual Responsibility." *The Monist*. 94 (3): 349–368.
- Hope, Simon. 2014. "Kantian Imperfect Duties and Modern Debates over Human Rights." *Journal of Political Philosophy*. 22 (4): 396–415.
- Horgan, Terry and Timmons, Mark. 2010. "Untying a Knot from the Inside Out: Reflections on the "Paradox" of Supererogation." *Social Philosophy and Policy* 27 (2): 29–63.
- Igneski, Violetta. 2006. "Perfect and Imperfect Duties to Aid." *Social Theory and Practice*. 32 (3): 439–466.
- Kagan, Shelly. 1989. *The Limits of Morality*. New York: Oxford University Press.
- Kagan, Shelly. 1991. "Replies to My Critics." *Philosophy and Phenomenological Research*. 51 (4): 919–928.
- Kagan, Shelly. 2011. "Do I Make a Difference?" *Philosophy and Public Affairs* 39 (2): 105–141.
- Kant, Immanuel. 1785/2012. *Groundwork of the Metaphysics of Morals*. Translated by M. Gregor and J. Timmermann. Cambridge: Cambridge University Press.
- Kant, Immanuel. 1797/2017. *The Metaphysics of Morals*. Translated by Mary Gregor. Cambridge: Cambridge University Press.
- Keefe, Rosanna. 2003. *Theories of Vagueness*. Cambridge: Cambridge University Press.
- Kingston, Ewan and Sinnott-Armstrong, Walter. 2018. "What's Wrong with Joyguzzling?" *Ethical Theory and Moral Practice*. 21: 169–186.

- Korsgaard, Christine. 1996. *Creating the Kingdom of Ends*. Cambridge: Cambridge University Press.
- Lea, David. 2004. "The Imperfect Nature of Corporate Responsibilities to Stakeholders." *Business Ethics Quarterly*. 14 (2): 201–217.
- McCarty, Richard. 1989. "The Limits of Kantian Duty, and Beyond." *American Philosophical Quarterly*. 26 (1): 43–52.
- McElwee, Brian. 2015. "What is Demandingness?" In van Ackeren, M. and Kühler, M. (eds.), *The Limits of Moral Obligation: Moral Demandingness and Ought Implies Can*, pp. 19–35. New York: Routledge.
- McNamara, Paul. 2011. "Supererogation, Inside and Out: Toward an Adequate Scheme for Common Sense-Morality." In M. Timmons (ed.) *Oxford Studies in Normative Ethics, Volume 1*, pp. 202–235. Oxford: Oxford University Press.
- Murphy, Liam. 2000. *Moral Demands in Nonideal Theory*. New York: Oxford University Press.
- Nefsky, Julia. 2011. "Consequentialism and the Problem of Collective Harm: A Reply to Kagan." *Philosophy and Public Affairs* 39 (4): 364–395.
- Nefsky, Julia. 2015. "Fairness, Participation, and the Real Problem of Collective Harm." In M. Timmons (ed.) *Oxford Studies in Normative Ethics, Volume 5*, pp. 245–271. Oxford: Oxford University Press.
- Nefsky, Julia. 2017. "How You Can Help, Without Making a Difference." *Philosophical Studies* 174 (11): 2743–2767.
- Nefsky, Julia. 2019. "Collective Harm and The Inefficacy Problem." *Philosophy Compass* 14 (4). <https://doi.org/10.1111/phc3.12587>.
- Nefsky, Julia. 2021. "Climate Change and Individual Obligations: A Dilemma for the Expected Utility Approach, and the Need for an Imperfect View." In Budolfson, Mark; McPherson, Tristram; Plunkett, David (eds.), *Philosophy and Climate Change*, pp. 201–221. Oxford: Oxford University Press.
- Noggle, Robert. 2009. "Give Till It Hurts? Beneficence, Imperfect Duties, and a Moderate Response to the Aid Question." *Journal of Social Philosophy*. 40 (1): 1–16.

- Ohreen, David and Petry, Roger. 2012. "Imperfect Duties and Corporate Philanthropy: A Kantian Approach." *Journal of Business Ethics*. 106 (3): 367–381.
- O'Neill, Onora. 1996. *Towards Justice and Virtue: A Constructive Account of Practical Reasoning*. Cambridge: Cambridge University Press.
- O'Neill, Onora. 2013. *Acting on Principle: An Essay on Kantian Ethics*. New York: Cambridge University Press.
- Otsuka, Michael. 1991. "The Paradox of Group Beneficence." *Philosophy and Public Affairs* 20 (2): 132–149.
- Parfit, Derek. 1984. *Reasons and Persons*. Oxford: Clarendon Press.
- Parfit, Derek. 2011. *On What Matters: Volume One*. Oxford: Oxford University Press.
- Pinkert, Felix. 2015. "What If I Cannot Make a Difference (and Know It)." *Ethics*. 125 (4): 971–998.
- Portmore, Douglas. 2011. *Commonsense Consequentialism: Wherein Morality Meets Rationality*. Oxford: Oxford University Press.
- Portmore, Douglas. 2018. "Maximalism and Moral Harmony." *Philosophy and Phenomenological Research*. (2): 318–341.
- Portmore, Douglas. Forthcoming. "Latitude, Supererogation, and Imperfect Duties." In D. Heyd (ed.) *Springer Handbook of Supererogation*. Springer.
- Pummer, Theron. 2016. "Whether and Where to Give." *Philosophy and Public Affairs* 44 (1): 77–95.
- Pummer, Theron. 2023. *The Rules of Rescue: Cost, Distance, and Effective Altruism*. New York: Oxford University Press.
- Raffman, Diana. 2012. "Indiscriminability and Phenomenal Continua." *Philosophical Perspectives*. 26 (1): 309–322.
- Rainbolt, George. 2000. "Perfect and Imperfect Obligations." *Philosophical Studies*. 98 (3): 233–256.
- Rawls, John. 1999. *A Theory of Justice*. Cambridge, Massachusetts: Harvard University Press.

- Regan, Donald. 2000. "Perceiving Imperceptible Harms (With Other Thoughts on Transitivity, Cumulative Effects, and Consequentialism)." In Michael Almeida (ed.), *Imperceptible Harms and Benefits*, pp. 49–71. Dordrecht: Kluwer.
- Sainsbury, Mark. 1988. *Paradoxes*. Cambridge: Cambridge University Press.
- Salam, Abdallah. 2014. "Perfect and Imperfect Rights, Duties and Obligations: from Hugo Grotius to Immanuel Kant." PhD thesis, University of Oxford.
- Scheffler, Samuel. 1992. *Human Morality*. New York: Oxford University Press.
- Schneewind, Jerome B. 1990. "The Misfortunes of Virtue." *Ethics*. 101 (1): 42–63.
- Schroeder, S. Andrew. 2014. "Imperfect Duties, Group Obligations, and Beneficence." *Journal of Moral Philosophy*. 11 (5): 557–584.
- Schumaker, Millard. 1992. *Sharing Without Reckoning: Imperfect Right and the Norms of Reciprocity*. Waterloo, Ontario: Wilfrid Laurier University Press.
- Segal, Aaron Eli. 2022. "Imperfect Duties and Good Will." PhD thesis, University of Pittsburgh.
- Shiffrin, Seana. 1991. "Moral Autonomy and Agent-Centred Options." *Analysis*. 51 (4): 244–254.
- Singer, Peter. 1972. "Famine, Affluence, and Morality." *Philosophy and Public Affairs* 1 (3): 229–243.
- Spiekermann, Kai. 2014. "Small Impacts and Imperceptible Effects: Causing Harm with Others." *Midwest Studies in Philosophy* 38 (1): 75–90.
- Statman, Daniel. 1996. "Who Needs Imperfect Duties?" *American Philosophical Quarterly*. 33 (2): 211–224.
- Sticker, Martin and van Ackeren, Marcel. 2018. "The Demandingness of Beneficence and Kant's System of Duties." *Social Theory and Practice*. 44(3): 405–436.
- Stohr, Karen. 2011. "Kantian Beneficence and the Problem of Obligatory Aid." *Journal of Moral Philosophy*. 8 (1): 45–67.
- Tenenbaum, Sergio and Raffman, Diana. 2012. "Vague Projects and the Puzzle of the Self-Torturer." *Ethics*. 123 (1): 86–112.
- Tiefensee, Christine. 2022. "Indeterminacy and Collective Harms." *Philosophical Studies* 179 (11): 3307–3324.

- Timmerman, Travis. 2015. "Sometimes There Is Nothing Wrong with Letting a Child Drown." *Analysis*. 75 (2): 204–212.
- Timmermann, Jens. 2005. "Good but Not Required? – Assessing the Demands of Kantian Ethics." *Journal of Moral Philosophy*. 2 (1): 9–27.
- Timmons, Mark. 2021. *Kant's Doctrine of Virtue*. New York: Oxford University Press.
- Urmson, J.O. 1958. "Saints and Heroes." In A.I. Melden (ed.) *Essays in Moral Philosophy*, pp. 198–216. Seattle: University of Washington Press.
- van Ackeren, Marcel. 2018. "How Morality Becomes Demanding: Cost vs Difficulty and Restriction." *International Journal of Philosophical Studies*. 26 (3): 315–334.
- Velleman, David J. 1997. "How to Share an Intention." *Philosophy and Phenomenological Research*. 57 (1): 29–50.
- Voorhoeve, Alex and Binmore, Ken. 2006. "Transitivity, the Sorites Paradox, and Similarity-Based Decision-Making." *Erkenntnis*. 64 (1): 101–114.
- Wallace, R. Jay. 2019. "Discretionary Moral Duties." In M. Timmons (ed.) *Oxford Studies in Normative Ethics, Volume 9*, pp. 50–72. Oxford: Oxford University Press.
- Wick, Warner. 1983. "Kant's Moral Philosophy." Introduction to *Ethical Philosophy: The Complete Texts of "Grounding of the Metaphysics of Morals" and "Metaphysical Elements of Virtue"*, pp. xi–lxii. [by Immanuel Kant]. Translated by James W. Ellington. Indianapolis: Hackett.