MORALITY AS UNIVERSAL LOVE:

IS LOVE ALL YOU NEED?

A Dissertation

presented to

the Faculty of the Graduate School

at the University of Missouri-Columbia

In Partial Fulfillment

of the Requirements for the Degree

Doctor of Philosophy

---

by

ALOK TIWARI

Dr. Peter Vallentyne, Dissertation Supervisor

JULY 2025

The undersigned, appointed by the dean of the Graduate School, have examined the dissertation entitled

MORALITY AS UNIVERSAL LOVE:

IS LOVE ALL YOU NEED?

presented by Alok Tiwari,

a candidate for the degree of Doctor of Philosophy,

and hereby certify that, in their opinion, it is worthy of acceptance.

_____

Professor Peter Vallentyne

_____

Professor Robert N. Johnson

_____

Professor Kenneth Boyce

_____

Professor Paul J. Litton

ACKNOWLEDGEMENTS

This dissertation has come to its completion only because of the generous help of many individuals who have supported me, professionally and personally, along my long journey of graduate studies in philosophy.

First of all, I would like to thank my dissertation committee. Peter Vallentyne has been enormously helpful throughout the dissertation writing process, and my time at Mizzou more generally. I have learnt how to write academic philosophy from him. I am especially grateful for his willingness to be generous with his time in retirement. I have also immensely benefited from Robert Johnson's insightful feedback on my work. I remain open to his suggestion that I may be a Kantian without realizing it. Kenneth Boyce has been a source of guidance and inspiration throughout my graduate career in more ways than one, and I am grateful for all his support through the years. I am also grateful to Paul Litton for making time in his busy schedule, and for providing probing feedback.

I would also like to thank all the faculty members, graduate students, and office staff, who have supported me during my graduate studies at Mizzou. In particular, I would like to thank Andrew Melnyk, Alex Radulescu, Andre Ariew, Marta Heckel, Paul Weirich, and Peter Markie, for their classes which have been vital in my growth as a student of philosophy. I would also like to thank Philip Robbins, M Radulescu, and Troy Hall, for helping me learn how to teach philosophy. I am also immensely grateful to all the graduate students at Mizzou, many of whom I consider dear friends. In particular, I would like to thank Ryan Kulesa, Gesiel da Silva, Nathaniel Stagg, Se Yong Bae, Xihe

Ouyang, Joel Vall Thomas, Christian Culak, Chuanke Wei, Jean Janasz, Haeseong Shin, Tianqin Ren, Davi Bastos, Naveed Rahman, Rajiv Gupta, Jacob Nofal, Nick Parker, Beck Apicello, Lauren Graf, Bryce Counts, Dennis Velez, Argon Gruber, Sukhvinder Shahi, Ritwik Agarwal, and Alek Willsey. Special thanks to Gesiel for suggesting the second half of the dissertation title.

I would also like to thank all the faculty members, graduate students, and office staff, at the University of Hyderabad who helped me begin my academic career in philosophy. In particular, I owe a deep debt of gratitude to Prajit Basu. His class on ancient Greek philosophy piqued my interest in philosophy, and it was only because of his encouragement that I considered coming to the USA to pursue a PhD. He has been a trusted mentor through all these years. I am also grateful to have had Kedar Patwary, Tias Maity, and Vijay Govind Nath, as my peers and friends.

I am also grateful to my friends in Missouri who have given me much-needed support throughout my graduate studies at Mizzou. I am especially grateful to have had Kaleb Watkins as a dear friend, along with Jennifer, Rowan, Piper, Oliver, and the extended family who have been like family to me. I am also grateful to members of Campus Lutheran Church for their hospitality. In particular, I would like to thank Bayliee Kulesa, Robert Meyer, Hope Kane, Beth Busseau, Jon Boldt, Cindy and Duane Craig, Aaron and Natalie Schmidt, Blake and Emilia Trickey, Mark and Kristy Kruse, and Pastors Kent Pierce and James Lanning.

Lastly, I owe an immense debt of gratitude to my parents and sisters for their constant support throughout my graduate career. In particular, none of my academic achievements would have been possible without my parents' unwavering support.

TABLE OF CONTENTS

MORALITY AS UNIVERSAL LOVE:

IS LOVE ALL YOU NEED?

Alok Tiwari

Dr. Peter Vallentyne, Dissertation Supervisor

ABSTRACT

In this dissertation, I investigate whether universal love can play a foundational role in moral theory. I argue that universal love is apt to play a foundational role in virtue ethical moral theory. I start by sketching a volitional account of love and use it to develop an account of universal love as the desire for the well-being of all beings with moral standing. I argue that universal love cannot play a foundational role in deontological or consequentialist moral theory, but it can play a foundational role in virtue ethical moral theory. I sketch a virtue ethical moral theory in which universal love is the sole cardinal moral virtue that grounds all the other moral virtues and defend the theory against some important objections.

1 Introduction

The concept of universal love plays a central role in many influential ethical traditions, including Buddhist and Christian ethics. In modern moral philosophy, on the other hand, the concept of love and especially universal love has received scant attention. Recently, however, there has been some interest among moral philosophers in exploring the relevance of love for moral theory. In this dissertation, taking the recent philosophical discussion as my starting point, I investigate the plausibility of explaining morality in terms of universal love.

In this introductory chapter, I first briefly introduce the concept of universal love as it appears in Buddhist and Christian ethics. I then briefly discuss recent philosophical work on the relationship between love and morality, in order to highlight some claims and issues raised in the literature. I then clarify some basic assumptions I make about morality as a normative domain throughout this dissertation. I conclude the chapter by providing a brief roadmap of what follows in the subsequent chapters.

1.1 Universal Love in Buddhist and Christian Ethics

Universal love emerges as a central concept in Buddhist and Christian ethics. Consider the following passages from the Buddhist Pāli Canon and the Christian New Testament.

"May all beings be happy and secure, may they be happy-minded.
Whatever beings there are… without exception be happy-minded.
Let none deceive another nor despise any person whatever in any place;
in anger or ill-will let them not wish any suffering to each other.
Just as a mother would protect her only child at the risk of her own life,
even so, let him cultivate a boundless heart towards all beings.
Let her thoughts of boundless lovingkindness pervade the whole world:
above, below and across, without obstruction, without any hatred, without any enmity."

Kho 9. Karaṇīya Mettā Sutta: 3-8

"But I say to you who are listening: Love your enemies; do good to those who hate you; bless those who curse you; pray for those who mistreat you."

Luke 6:27-28

The Pāli Canon prescribes the ideal of *mettā*, translated as boundless lovingkindness above, while the New Testament prescribes the ideal of *agapē*, roughly unconditional love. The two ideals are not identical, but they do have significant overlap in their content. In particular, both *mettā* and *agapē* involve a desire for others' well-being and are unconditional, in that there are no conditions to be met for someone to be an object of this desire. What is common to both is a generic notion of what we might call *universal love* as an ideal. Let us unpack this notion.

Part of what is being prescribed in the Buddhist and Christian scriptures is an *affective disposition*. The *mettā sutta*, for instance, forbids *despising* any living being or *wishing* that they suffer. Jesus, according to Luke's account, offers a similar prescription to *bless* others and *pray* for them — both of which express pro-attitudes towards other's well-being. Further, part of the prescription is also *practical*, that is having to do with *action* that directly affects the well-being of others.[1] Jesus gives a direct prescription to do good to others, which plausibly means acting in a way that improves others' well-being. The *mettā sutta* likewise prohibits any deception. In addition to an affective disposition of taking pleasure in others' well-being, then, there is also a practical disposition to improve others' well-being through one's actions. Together, these affective

[1] Of course, praying for someone or blessing them is also an *action* but it is not *practical* in the relevant sense as it does not *directly* affect their well-being. While praying for someone might lead to divine intervention and thus *indirectly* affect their well-being, the effect on their well-being is not a direct result of one's action.

and practical dispositions might be thought to be constitutive of the *desire* for the well-being of others that is a part of the ideal of universal love.

The striking feature of the Buddhist ideal of *mettā* and the Christian ideal of *agapē* is their *unconditional* nature. The *mettā sutta* is emphatic in its assertion of the exceptionless and boundless nature of *mettā* — it does not discriminate between the weak and the strong, those living near and far away, those already born and those not yet born. Jesus similarly emphasizes the radical nature of *agapē* — one is to love not just those who love oneself but also those that hate oneself, one is to do good not just to those who do good to oneself but also to those who harm oneself. The result is an *ideal* of how one ought to relate to others. In particular, the *unconditional* nature of universal love stands in sharp contrast with other ways of relating to other beings that is *contingent*, say on whether one's concern for the other is reciprocated.

We may abstract from *mettā* and *agapē*, then, and say that universal love is an unconditional desire for the well-being of others. Of course, the notion needs to be specified in much greater detail before it can be used for any explanatory purposes. An obvious interpretative question, for instance, is the scope of universal love. The scope of the Buddhist ideal of *mettā* explicitly includes all sentient beings. The Christian ideal of *agapē* is not explicit about its scope, but it is typically interpreted as applying to all and only human beings. I shall return to these interpretive issues in a later chapter, for now I hope to have said enough to give the reader some grasp of the concept of universal love.

1.2 Moral Philosophers on Love and Morality

The concept of love and especially the concept of universal love has been of relatively little interest among modern moral philosophers. There has, however, been some interest

in recent moral philosophy on the relationship between love and morality. I briefly survey this recent work to highlight some claims and issues raised in the literature.

David Velleman argues that love is a moral emotion that makes one aware of the intrinsic value of other rational beings (Velleman 1999, 2008). On his view, respect is the *required* minimal response to someone's rational capacity of valuation, while love is an *optional* maximal response to the same capacity. The idea is that respect for the rational nature of others restrains one's self-interested tendency to use them as mere means. *Love* goes further and makes one intensely aware of the intrinsic value of the other's rational capacities. Love, then, makes one emotionally vulnerable to other rational beings. Love, like respect, allows an agent to recognize the intrinsic value of other rational beings and is therefore properly thought of as a moral emotion.

Kieran Setiya similarly argues that it is rational for a human moral agent to love any other human being (Setiya 2014, 2022). That is, the fact someone is a human being is enough to justify loving them. An implication of this fact for moral philosophy, that he explicitly draws out, is that in rescue cases one is *not* required to save the greater number. Roughly, his argument is that because one is justified in loving *any* human being and love justifies acting with partiality towards one is beloved, one is justified in loving the fewer number and therefore rescuing them rather than the larger number. More generally, an important implication of his views is that if morality is based on love, morality does not require impartial aggregation of claims or well-being.

Talbot Brewer argues that love reveals to us the irreplicable value of fellow human beings (Brewer 2018). He focuses not on universal love, but on the love for particular individuals with whom one is in a loving relationship. His focus is thus on

philia, the love for someone in particular with whom one shares a loving relationship, rather than the unconditional love of *agapē*, which is not restricted to those with whom one shares a loving relationship. He argues that relationships with other human beings gives us access to their irreplaceable value, which we would not otherwise have. Further, the access to the irreplaceable value of others also helps us grasp the fact other human beings are owed certain forms of treatment precisely in virtue of having this irreplaceable value.

Jay Wallace makes a related argument for the claim that there are *sui generis* duties of love, that is, duties we owe to others in virtue of being in loving relationships with them (Wallace 2012). He argues for a non-reductionist view about these special duties, against reductionist views that try to reduce these special duties to more general moral principles. The upshot for moral philosophy is that duties of love give us paradigmatic instances of *relational* obligations — obligations owed to someone in particular. He argues that obligations in general are relational, so it would be wrongheaded to try to explain duties of love, a paradigmatic instance of relational obligations, in terms of non-relational moral principles. He suggests that rather than trying to reduce duties of love to some general moral principles by appealing to impersonal moral values, it is more promising to try to generalize from special duties of love to moral obligations in general.

This has so far been a very rough sketch of the recent philosophical literature on love and morality. Rather than explaining or assessing the arguments in detail, I want to highlight a few points from this brief review. First, the fact that moral philosophers of rather different philosophical orientations find the concept of love illuminating for

explaining moral phenomena gives us reason to explore the plausibility of morality as universal love. In particular, Velleman finds the concept of love to have explanatory value within a broadly Kantian framework, while Brewer locates the explanatory value of love within a broadly Aristotelian framework. Whatever the merits of their particular arguments, their attempts give us reason to think that love might play an important role in moral theory.

Second, most of the philosophical work has been focused on other human beings as the object of love. Velleman focuses on rational beings as objects of love, rather than human beings per say. Setiya, on the other hand, explicitly argues that non-rational human beings are our moral equals and therefore equally deserving of love (Setiya 2023). Whether all human beings have equal moral value is an interesting issue, though I am not convinced that we have good reason to restrict love to human or rational nature. I think that love could have as its object, not just a fellow human or rational being, but also, a fellow animal creature. Of course, whether or not love for animals is possible will ultimately depend on one's account of love.

Third, even if Wallace is right that there are sui generis duties *of* love, for love to be foundational to morality there must also be a duty *to* love. Or at least, there must be a duty to *cultivate* love. For, duties of love, as Wallace understands them, are duties owed to others insofar as one loves them. If all moral duties, including duties of justice, are ultimately duties of love, these duties would arise only insofar as one loves beings with moral standing. It is possible, however, that an agent may owe some duties to someone that she does not love. To put it more abstractly, moral duties have a *necessity* that could not be explained if moral duties are *contingent* on an agent's love for beings with moral

standing. If love is to play a fundamental role in explaining moral duties, there must be a moral duty to love beings with moral standing or at least a moral duty to cultivate such love. Moreover, one might plausibly disagree with Wallace those duties, relational or otherwise, are at the heart of morality in the first place. So even if there are duties of love, it may be mistaken to think of them as central to morality.

Lastly, Brewer's epistemological claim that *philia*, love for particular individuals with whom one is in a loving relationship, gives us unique access to the irreplaceable value of others is compatible with the metaphysical claim that moral obligation is grounded in *agapē*, that is universal love. For instance, it is plausible that one comes to recognize the value of animal life only by having a pet. But that does not entail that the value of animal life is grounded in such relationships. Rather, one's relationship with one's pet might give one epistemic access to the intrinsic value of animal life. Importantly, any moral duties to one's pet might be grounded ultimately in the intrinsic value of animal life. More generally, then, it may be the case that moral obligations are grounded in universal love, even though we come to recognize these duties only by being in loving relationships with particular others.

1.3 Moral Normativity

I have so far briefly introduced the role of universal love in Buddhist and Christian ethics, as well as the recent philosophical discussion on the relationship between love and morality. I now want to lay the groundwork for what follows by clarifying my assumptions about morality as a normative domain.

I take morality to be one normative standpoint, among others. A normative standpoint can be understood, following Dale Dorsey (2016), in *functional* terms. A

normative standpoint takes as inputs certain *facts* about an action, belief, desire etc., and generates a corresponding *evaluative output* of "good," "bad," "permissible," "impermissible" etc. Morality, then, is one normative standpoint alongside others, like prudence, which can be used to evaluate actions, beliefs, or desires.

Further, any given normative standpoint generates standpoint-relative *reasons* of different weights or strengths — stronger or weaker considerations for or against what is being evaluated. When there is *decisive* standpoint-relative reason to *x*, there is a standpoint-relative *requirement* to *x*. As I understand it, a standpoint-relative requirement just is a standpoint-relative duty. A *moral duty* to do something, then, arises when there is *decisive moral reason* to do something. In other words, an agent has a moral duty to do something when she has decisive reason to do it from the moral point of view. In general, whether the standpoint-relative reasons are decisive would depend on the particular standpoint in question. On a maximizing conception of moral requirements, for instance, one is morally required to do what one has *most* moral reason to do.

Also, a general assumption I make is that an agent has the faculty of *practical reason* that enables her to answer the question of what she should *do* by weighing all the relevant kinds of reasons for action in practical deliberation. That is, practical reason enables an agent to take up a *comprehensive* point of view which takes into consideration all the relevant normative standpoints to answer the question of what she should do, all-things-considered. For instance, consider the possible action of wading into a shallow pond to save a drowning child. The fact that doing so will save the child's life generates a *moral requirement* to perform the action. At the same time, the fact that performing the action would wet one's shoes is *prudentially bad*. The agent's practical reason enables

her to recognize that she has not just decisive moral reason to perform the action but that she has decisive all-things-considered reasons to perform the action.

Now, if morality is one normative standpoint among others, what distinguishes the moral standpoint from other normative standpoints? In what follows, I rely on a pre-theoretic understanding of the moral domain. Nonetheless, it will be useful, though no doubt controversial, to lay out three features that I take to be distinctive of the moral domain — accountability, impartiality, and concern for welfare. None of these individual features is unique to morality, but collectively they are arguably distinctive of morality. Let me briefly explain each feature.

First, the moral standpoint is one from which moral agents hold one another accountable for complying with moral requirements. For instance, a moral agent who fails to comply with a moral requirement without sufficient justification may justifiably be blamed by other moral agents. Consider again the drowning child case — if the agent does not wade into the pond, she may justifiably be blamed for her failure to rescue the drowning child. Accountability is not *unique* to morality since the law holds agents accountable for complying with legal requirements. Accountability does, however, does help us to distinguish morality from some other normative domains. For instance, arguably agents cannot justifiably hold one another accountable for complying with prudential requirements. Say an agent has a prudential duty not to smoke a pack of cigarette every day but does so anyway. Arguably, the agent's action does not warrant blame. The agent's loved ones may permissibly counsel her that it would be in her best interest to stop smoking, and they may even be required to do so. But they would not be warranted in *blaming* her for failing to heed their counsel. One relevant difference

between the prudential case and the moral case is that in the moral case what the agent does has direct significance for others, while the prudential case is the agent's business alone.

Second, the moral standpoint is impartial. Some normative domains, like the norms of familial relationships and friendship are inherently partial and take into consideration only the agent's family and friends. Moral norms, on the other hand, are impartial insofar as they take into consideration *all* the relevant parties who stand to be affected by, say, one's actions. In principle, the relevant extension of "all" could be all human beings, all persons, all sentient beings, etc. depending on one's substantive views on moral standing. The important point, though, is that morality is not a partial domain concerned only with some, but an impartial domain concerned with all, however exactly one fills in the details. I should also clarify that impartiality does not entail *equality*. For instance, to say that pigs count from the moral point of view does not entail that pigs count morally *as much as* human beings. Moral impartiality might require taking pigs into consideration but that is compatible with the claim that human beings matter more than pigs from the moral point of view.

Third, the moral standpoint is concerned with the well-being of the concerned parties. I have a broad conception of well-being in mind here — anything that is *good for* someone contributes to their well-being. For instance, exercising one's autonomous capacities is plausibly good for a rational being. Similarly, avoiding pain is good for a sentient being. Morality is concerned with well-being, unlike some other normative domains. Etiquette, for instance, is concerned not with others' well-being but with social norms of politeness. Wearing sports shoes to a formal dinner may violate a norm of

etiquette without harming anyone's well-being. Violations of moral norms, on the other hand, paradigmatically negatively affect others' well-being. More generally, morality is concerned with the well-being of others while etiquette is not. This third feature might perhaps be the most controversial of the three features, but it arguably helps us to distinguish morality from other normative domains.

1.4 Aim and Structure of the Dissertation

In this dissertation I shall investigate the plausibility of explaining morality in terms of universal love. Minimally, a plausible account of morality will provide an adequate explanation of the moral properties of actions as well as the moral character of agents. The central question of this dissertation is: Can universal love play a foundational role in a plausible account of morality? In short, I shall answer this question in the affirmative. In particular, I shall argue that universal love is well suited to play a central role in a virtue ethical moral theory.

Here is how I shall proceed in what follows. I start by sketching an account of love in chapter 2. Briefly, on my view, to love someone is to desire the promotion of their well-being for their own sake. Using this account of love, I develop an account of *universal* love in chapter 3. On my account, universal love is the desire for the promotion of the well-being of all beings with moral standing (that one veridically believes to exist). That is, universal love is the species of love that has as its object all beings with moral standing (that one veridically believes to exist). Having developed an account of universal love, I address the central issue of this dissertation — the role that universal love can play in moral theory — in chapters 4 and 5.

In chapter 4, I argue that universal love can play a foundational role in a virtue ethical moral theory, but not in a deontological or consequentialist moral theory. In particular, I argue that universal love cannot be the object of a moral requirement because there cannot be a *duty* to love. As a result, universal love cannot play a foundational role in a deontological moral theory. I then suggest that the disposition to love universally is a moral virtue. That is, the disposition to love all beings with moral standing is a morally excellent disposition. In chapter 5, I sketch a moral theory in which the disposition to love universally is the sole cardinal moral virtue that grounds all the other moral virtues but is not itself grounded in any other moral virtue. In particular, I argue that the commonly recognized moral virtues, like benevolence and justice, are different aspects of the disposition to love universally. Lastly, in chapter 6, I defend this moral theory by considering and answering some important objections.

2 A Volitional Account of Love

In ordinary discourse we speak of 'love' for all kinds of things — foods, activities, places, colors, seasons etc. Even if this usage of "love" picks out some broad notion of love, it is doubtful that this notion is relevant for moral theory. Rather, the sort of love that might plausibly be relevant for moral theory is love for one's spouse, friends, children, or pets. What is distinctive of this latter sort of love is that it is love for *someone*, rather than love of *something*.[2] It is this notion of love for someone that I have in mind in what follows. Also, to be clear, I assume that love for someone is a genus with many species, including romantic love, friendship love, and parental love.

In this chapter, I provide an account of love that I will use throughout this dissertation. In the first section I provide a sketch of my account. In the second section, I clarify my account by contrasting it with two competing accounts of love.

2.1 Proposal: An Account of Love

On my view, to love someone is to desire the promotion of their well-being for their own sake. I understand desires as *pro tanto* dispositions to act: An agent *A* desires *x* if and only if *A* is *pro tanto* disposed to exercise her agency to bring about *x*. To love someone, then, is to have the *pro tanto* disposition to promote their well-being through the exercise of one's agency. This includes the *pro tanto* dispositions to (i) perform actions that one believes will improve their well-being, and (ii) refrain from performing actions that one believes will reduce their well-being.

[2] Many accounts of love in the literature focus on love of persons (e.g., Velleman) or human beings (e.g., Setiya). The notion of love that I am interested in is restricted neither to love of persons or human beings but includes love of non-persons like infants and non-humans like pets. I take love of persons to be a species of this notion of love. I elaborate on this point in the next chapter.

In the rest of this section, I elaborate on the two central elements of this account — desire and well-being — and explain how they together constitute love on my view.

2.1.1 Desire

I take a desire for something to be a *pro tanto* disposition to exercise one's agency in order to bring it about.[3] Here, by "to exercise one's agency" I mean an exercise of one's will with the intention to bring about a certain result. That is, an agent exercises her agency when she *does* something that she *believes* will bring about a result that she *intends* to bring about. This means that an agent does *not* exercise her agency when she acts on a mere reflex or instinct. For instance, if an agent scratches her forehead on a mere reflex, she does not exercise her agency because she does not intend to scratch her forehead. I treat "desire" as synonymous with "want". So, another way to put the same point is that in this case the agent does something without *wanting* to do it. When an agent acts on a desire, however, she does something because she wants to.

Consider a mundane desire like a desire to eat ice-cream. Having the desire to eat ice-cream, as I understand it, is the state of being *pro tanto* disposed to exercise one's agency to bring it about that one eats ice-cream. That is, to have the desire is to be *pro tanto* disposed to do what one believes will bring it about that one eats ice-cream. For instance, the agent would be *pro tanto* disposed to visit the local ice-cream store or to stock up her freezer with ice-cream from the grocery store. Of course, an agent may have other competing desires, for instance a desire to lose weight or save money, which will

[3] This is an action-based theory of desire. This is no doubt a controversial theory of desire for an overview of the different theories of desire, see the Stanford Encyclopedia of Philosophy entry on Desire.

determine the frequency with which she acts on her *pro tanto* disposition to eat ice-cream.

Desires, then, are motivations for action that help explain intentional actions. That is, an agent A's desire for *x* gives A *pro tanto* motivation to act to bring about *x*. Of course, the agent may have other desires which might make it the case that she does not have *conclusive* motivation to act to bring about *x* in any given circumstance. Whether she acts on a particular desire in any given case depends on the competing desires she could bring about in that particular case and their relative strengths. An agent *A*'s desire for *x* is stronger than her desire for *y* if and only if *A* has stronger motivation to act to bring about *x* than she does to bring out about *y*.

On my view, then, to desire something is to want it to be the case. To want something to be the case is to be *pro tanto* disposed to bring it about.

2.1.2 Well-Being

The term "well-being" is used in a narrow as well as a broad sense. In the narrow sense of the term, someone's well-being is a matter of how *well* their life goes for them. That is, in the narrow sense of the term, someone's well-being refers to whatever is intrinsically, that is non-instrumentally, *good for* them. Pleasure, for instance, might constitute someone's well-being in the narrow sense. In the broad sense of the term, on the other hand, someone's well-being refers to how their life goes for them *overall*, which includes the intrinsic goods as well as the intrinsic bads in their life. In the broad sense of the term, someone's well-being might be constituted by the balance of pleasure over pain. I use

"well-being" in the latter broad sense. That is, I use "well-being" to refer to the overall prudential value of someone's life.

The *elements* of someone's well-being contribute, positively or negatively, to their well-being. The *positive elements of well-being* contribute to making someone's life go better. The presence of a positive element of well-being in one's life makes it the case that one's life is better than it would have been in its absence. Similarly, the *negative elements of well-being* are things that contribute to making someone's life go worse. The presence of a negative element of well-being makes it the case that one's life is worse than it would have been in its absence. Someone's overall level of well-being will depend on the balance of the positive and negative elements of well-being in their life. To be clear, I take no stance on how different elements of well-being are to be balanced. For all we know, this might turn out to be a rather difficult balancing act.

Further, on my usage of "well-being", the subject of well-being could be a person or an animal. For instance, we could ask of a dog, as much as of a person, how their life goes for them. Of course, the standards that inform our answers in the two cases will be different, but the point is that we could ask the same question in each case — how does their life go for them?

Lastly, I assume that the elements of well-being are objective and plural. I assume, in other words, the so-called objective-list theory of well-being. I assume that there are multiple elements of well-being that contribute to the well-being of any being of that kind. I do not assume any particular exhaustive list of the elements of well-being, though I do make some relatively uncontroversial assumptions about some elements of well-being. In particular, I assume that pleasure, autonomy, knowledge, achievement, and

friendship are positive elements of well-being for human persons. I also assume that pain is a negative element of well-being of well-being for human persons. In the case of humans who are sentient but not yet rational, like infants, as well as sentient non-human animals, I assume that pleasure and pain are elements of their well-being. Once again, I am only making relatively uncontroversial assumptions about some elements of well-being, leaving open the possibility that there may be other elements not listed here.

2.1.3 Love

Having clarified what I mean by "desire" and "well-being" I now want to elaborate on how these two elements come together constitute love on my account. On my view, to love someone is to desire the promotion of their well-being for their own sake. Let me start by clarifying the first half of this account: What is it to desire the promotion of someone's well-being?

To begin with, well-being as an object of desire is quite distinct from ordinary objects of desire like eating ice-cream. Roughly, the object of the desire to eat ice-cream is a particular state of affairs — that of eating ice-cream. However, not all objects of desire are for particular states of affairs. Consider the desire for wealth. Typically, if one desires wealth one desires to have *more* wealth. So understood, the object of one's desire is a continuous variable rather than a particular state of affairs. The desire consists in the disposition to bring it about that one has *more of something* rather than a disposition to bring it about that one has *something*. Well-being as an object of desire is like wealth rather than eating ice-cream. That is, to desire someone's well-being is not to desire any particular state of affairs, but rather to desire the promotion of their well-being.

To desire the promotion of someone's well-being, in turn, is to have the *pro tanto* disposition to (i) perform actions that one believes will improve their well-being, and (ii) refrain from performing actions that one believes will reduce their well-being. This is because either forms of exercise of one's agency promote the beloved's well-being. For instance, if *A* loves *B*, *A* will be *pro tanto* disposed to perform actions that she believes will result in B's experiencing pleasure which will thereby increase *B*'s well-being. Similarly, *A* will be *pro tanto* disposed to refrain from performing actions that she believes will result in *B*'s experiencing pain which will thereby decrease *B*'s well-being. To be clear, the fact that an action *X* will cause *B* pain is only a *pro tanto* reason to refrain from performing *X*. It may be that in some cases *A* has overriding reason to perform *X*, say because performing *X* will save *B* from experiencing much greater pain in the future. I discuss these trade-offs in more detail in chapter 5.

What is it to desire the promotion of someone's well-being *for their own sake*? Let us start by considering the alternative. One could desire the promotion of someone's well-being for one's own sake. For instance, consider Siddhartha, a child, and Maya, his mother. Maya could desire the promotion of Siddhartha's well-being because it makes her feel good about herself as a parent. In this case, while Maya desires the promotion of Siddhartha's well-being, she does not *love* him because she desires the promotion of his well-being for the wrong reasons. For Maya's desire to constitute love, she must have an *intrinsic* desire for the promotion of Siddhartha's well-being. That is, if Maya loves Siddhartha, then his well-being motivates her to desire to promote it, rather than any perceived benefit to herself.

Desiring someone's well-being for merely instrumental reasons, then, is *not* love. Love consists rather in an *intrinsic* desire for someone's well-being.[4]

2.2 Alternative Accounts of Love

I have so far sketched a broadly *volitional* account of love. To help clarify my account, I now want to briefly contrast my account with competing *affective* and *cognitive* accounts of love. To be clear, my aim is not to provide an exhaustive survey of the literature, nor is my aim to argue for the superiority of my account over these alternative accounts. Rather, my aim is to briefly discuss these alternative accounts to help clarify what my account of love is *not*.

The volitional account of love is nicely summed up by Harry Frankfurt: "Of course, love does ordinarily involve various strong feelings and beliefs that express, reveal, and support it. The heart of love, however, is neither affective nor cognitive. It is volitional." (Frankfurt 1999) Love, on his view, is fundamentally about patterns of *willing* rather than *feeling* or *valuing*. Here is the basic idea. On the affective account love consists in appropriate affective dispositions concerning the beloved, while on the cognitive account love is an attitude of valuing someone that in turn leads to affective dispositions. On the volitional account, on the other hand, love consists in volitional dispositions which in turn leads to affective as well as cognitive dispositions. I shall now briefly discuss each of these alternative accounts and contrast them with my volitional account.

[4] This is not to deny that one could not have also have a prudential desire for the promotion of someone's well-being. Typically, the flourishing those one loves also improves one's own well-being. So, one desire the promotion of someone's well-being for their own sake as well as one's own sake.

2.2.1 The Affective Account of Love: Love as Feeling

On the affective account of love, to love someone is to have certain affective dispositions regarding them. Different affective accounts of love differ in their specification of the relevant affective dispositions. For instance, Neera Badhwar suggests that the disposition to feel "pleasure in (the thought of) the existence and well-being of the loved individual… and pain in (the thought of) his non-existence or ill-being, is central to the emotion of love." (Badhwar 2003:43-44) On her account, these dispositions are manifested in the paradigmatic cases but need not be manifested under *all* circumstances. For instance, if one's child has a debilitating terminal illness, it is consistent with love for one to *not* take pleasure in the child's existence. The idea, however, is that in the absence of such tragic circumstances love for one's child would manifest itself as pleasure in her existence.

Paul Helm similarly provides an affective account of love in terms of *person-focused emotions* (Helm 2009). On his view, person-focused emotions are emotions like pride and shame which have *particular persons* as their object. What is distinctive of these emotional responses is that they are rationally determined by the values constitutive of the identity of persons. The idea is that the identity of a person is constituted by *her* values. To care about a person, then, is to care about whether or not she lives up to the values that constitute her identity. This, in turn, entails that to care about a person is to be rationally committed to feeling person-focused emotions that take her as the object of one's emotional attitudes. For instance, it is a rational response to feel proud of oneself when one successfully completes a long and demanding project. Similarly, it is a rational response to feel ashamed when one gives up on the project mid-way. On Helm's account,

"… love just is a distinctive kind of affectionate, identificatory commitment to another: the kind of commitment that emerges from a rational pattern of person-focused emotions." (Helm 2009:52) For instance, one might be disposed to feel ashamed of a friend's actions when she fails to live up to her ideals, just as one might be disposed feel ashamed of oneself when one fails to live up to one's own ideals. This is meant to be a non-reductive account of love in terms of person-focused emotions.

Both Badhwar's and Helm's accounts, then, have in common the idea that to love someone is to be *emotionally vulnerable* to them. Where the accounts differ is in their specification of the relevant emotional dispositions — feelings of pleasure and pain, or intimate person-focused emotions of pride and shame.

Now, love does typically make one emotionally vulnerable to the beloved. On my view, however, the emotional vulnerability associated with love is not *constitutive* of love. Rather, the emotional vulnerability is the *result* of the volitional dispositions constitutive of love. Typically, the satisfaction of one's desires leads to a feeling of fulfillment while a frustration of one's desires produces a feeling of disappointment. These emotional dispositions associated with desires can be explained as *consequences* of the volitional dispositions constitutive of desires. The realization of something one desires leads to fulfillment while non-realization leads to disappointment that derives from the failure to realize one's desires. These feelings of fulfillment and disappointment must be distinguished from any pleasure or pain that is intrinsic to the object of one's desire. For instance, eating ice-cream is intrinsically pleasurable but that is distinct from the feeling of satisfaction that results from the realization of one's desire to eat ice-cream.

Similarly, the emotional dispositions associated with love can be explained as consequences of one's desire for the promotion of the beloved's well-being.

2.2.2 The Cognitive Account of Love: Love as Valuing

On the cognitive account of love, to love someone is to *value* them in some appropriate way. Contra the affective account, on the cognitive account what is central to love is the attitude of valuing someone, which in turn gives rise to the emotional dispositions characteristic of love. I briefly discuss two cognitive accounts of love.

On David Velleman's Kantian account, love of someone is the *awareness* of the inherent value of their *rational* nature (Velleman 1999). He argues that love is essentially an attitude toward the beloved as an end in herself, rather than an attitude toward any *outcome* concerning the beloved, like her well-being. On his view, love for someone is "… the awareness of a value inhering in its object… and an arresting awareness of that value." (Velleman 1999:360) In particular, he argues that any rational being warrants this awareness of the value of her rational nature. The idea is that a rational being's capacity to *value* makes her worthy not only of respect, but also of love. Respect is the *required* minimum response, while love is the *optional* maximum response to someone's rational capacity of valuation.

On Velleman's view, what is distinctive of the response of love is that "Love disarms our emotional defenses; it makes us vulnerable to the other." (Velleman 1999:361) The idea is that one can *respect* another rational being without being *emotionally vulnerable* to them. For instance, one can act to help someone in dire need purely out of respect without being emotionally distressed by their dire circumstances.

But if one does the same action out of love, one acts from an awareness of the value in the person. Presumably, one cannot be aware of this value without it leading one to be emotionally vulnerable to the other.

To be clear, on Velleman's view, emotional vulnerability is the *result* of love rather than being *constitutive* of love. On his view, love is fundamentally an attitude of valuing someone's rational nature which in turn makes one emotionally vulnerable to them as a result of one's awareness of their value, qua rational being.

Niko Kolodny presents an alternative cognitive account of love. On his account, love is valuing one's *relationship* to the beloved, rather than valuing the beloved herself (Kolodny 2003). On Velleman's account, love is an attitude of valuing someone for their *nonrelational* property of having the rational capacity of valuation. Against Velleman's account, Kolodny raises the objection that it succumbs to the nonsubstitutability objection. Roughly, the objection is that Velleman's account fails to explain the irreplaceability of the beloved because given his account one has just as much reason to, say, love one's spouse as any other person. But, the argument goes, one's reason to love one's spouse is a reason to love *her* rather than any other person.

On Kolodny's account, on the other hand, love is an attitude valuing someone for their *relational* property of being, for instance, one's child, husband, or sister. On his view, the reason that warrants love for someone is not the value inherent in their rational nature, or any other intrinsic nonrelational value. Rather, the reason that warrants love for someone is one's ongoing relationship with them. That is, what warrants love for one's spouse is precisely the fact she is one's spouse, rather than facts about her intrinsic

properties. In other words, it is relationships that are the source of reasons for love rather than individual persons.

Both Velleman's and Kolodny's accounts, then, have in common the idea that to love someone is to *value* them. Where the accounts differ is in their specification of the relevant object of one's valuational attitude — the person's intrinsic rational capacities or one's ongoing relationship with her.

Now, it is true that love is typically accompanied by such valuational attitudes. Such valuational attitudes give rise to an expression of one's love in art and poetry. On my view, however, the valuational attitudes associated with love are not constitutive of love itself. Rather, the valuational attitudes are *results* of the volitional dispositions constitutive of love. One's desire for something often, but not always, leads one to judge the object of one's desire to be good. Consider one's desire for ice-cream. If one were to reflect on one's desire and endorse one's desire for ice-cream, one would judge that ice-cream is good. But this need not always be so. For, it is possible that one has a desire to smoke but on reflection one does not endorse one's desire to smoke. That is, one may have a *pro tanto* disposition to smoke without valuing smoking. It is consistent with desiring to smoke that one judges smoking to be bad. Similarly, reflecting on one's intrinsic desire for the promotion of someone's well-being typically leads one to endorse the desire and value the beloved as well as one's relationship with her. But this need not always be the case. It is possible that on reflection, one may not endorse one's intrinsic desire for the promotion of someone's well-being. That is, one may love someone without valuing them or one's relationship with them. This could happen, for instance, in

abusive relationships where one loves someone, but neither values them nor one's relationship with them.

To sum up the chapter: To love someone is to desire the promotion of their well-being for their own sake. To desire the promotion of someone's well-being is to have the *pro tanto* dispositions to (i) perform actions that one believes will improve their well-being, and (ii) refrain from performing actions that one believes will reduce their well-being. These volitional dispositions, in turn, typically make one emotionally vulnerable to the beloved. They also typically lead one to value the beloved and one's relationship with them.

3. An Account of Universal Love

In chapter 1, I briefly introduced the notion of universal love as it appears in Buddhist and Christian ethics. Common to the Buddhist notion of *mettā* and the Christian notion of *agapē* is what we might call universal love — an unconditional desire for the well-being of others. Taking my cue from these traditions, I am investigating the possibility that the notion of universal love might play a foundation role in moral theory. To do so, I need to first clarify the notion of universal love. That is the task of this chapter. Having clarified the notion of universal love, in subsequent chapters I will investigate the role that universal love can play in moral theory.

I begin by sketching an account of universal love using the volitional account of love sketched in the previous chapter. Briefly, universal love is the desire for the promotion of the well-being of all beings with moral standing (that one veridically believes to exist). I then clarify my account by contrasting it with Christine Swanton's account of universal love.

3.1 Proposal: An Account of Universal Love

On my account, universal love is the desire for the promotion of the well-being of all beings with moral standing (that one veridically believes to exist). That is, universal love is the species of love that has as its object all beings with moral standing (that one veridically believes to exist). To manifest universal love, then, is to have the *pro tanto* dispositions to (i) perform actions that one believes will improve the well-being of all beings with moral standing (that one veridically believes to exist), and (ii) refrain from performing actions that one believes will reduce their well-being.

In the remainder of this section, I motivate and clarify this account. I first address the issue of whether there are normative reasons for universal love. I then clarify the scope of universal love. On my view there are normative reasons for universal love which also determine its scope. In particular, I suggest that the scope of universal love is all beings with moral standing (that one veridically believes to exist). I conclude with a discussion of the impartial and unconditional nature of universal love on my account.

3.1.1 The Reasons for Universal Love

Whether there are normative reasons for love in general is a controversial issue in the philosophical literature on love. The issue is whether there are any facts about someone that justify loving them. On Harry Frankfurt's view, for instance, there are no normative reasons for love. On David Velleman's view, on the other hand, the fact that someone is *rational* justifies, but does not require, love for them. Similarly, Kieran Setiya's view is that the fact that someone is *human* justifies, but does not require, love for them.

I agree with Velleman and Setiya that there are normative reasons for love in general. The relevance of this claim for our purposes is that if there are reasons for love in general, there are reasons for universal love as well. Accordingly, I now motivate the claim, though I do not try give a thorough defense of it.

Let me start by noting two sorts of cases where there are normative reasons for love. First, there are cases where an agent has normative reasons *not* to love someone. Consider Laila, a woman in an abusive marriage. Her husband has been abusing her physically as well as emotionally for a few years. Nonetheless, she continues to love her

abusive husband. In this case, Laila has reason not to continue loving her husband.[5] One can imagine a concerned friend asking her: "Why do you still love him?" And the question would be perfectly justified. Her friend's question reflects the fact that Laila's continued love for her husband is *bad*, if not *wrong*, given the abusive nature of the marriage. This is a case where an agent has reason to stop loving someone. One can similarly think of cases where one has reason not to start loving someone. For instance, one has reason not to become friends with someone who supports genocide.[6]

Second, there are cases where an agent has normative reasons to love someone. For instance, the fact that someone is one's child gives one reason to love them. It is *good*, if not *obligatory*, to love one's child. This is reflected in the fact that a parent who does not love their child can justifiably be asked: "How could you not love your own daughter?" To be clear, in some contexts the speaker might be expressing her puzzlement about how it is psychologically *possible* for a parent to not love their child. What I have in mind, though, is a speaker who is trying to get the unloving parent to reflect on the fact that they *should* love their child.

Now, these cases only show that there are normative reasons for love in *some* cases. Here is further motivation for thinking that there are normative reasons for love in *all* cases. On my account, to love someone is to desire the promotion of their well-being for their own sake. It follows that the fact that an entity is not a subject of well-being gives one normative reason *not* to love them on pain of incoherence. For, if one were to

[5] To be clear, I don't think Laila has reason to stop loving her husband altogether. On my view, she has reason to stop loving him qua her husband, but she has reason to continue loving him qua a human person.

[6] Once again, on my view, one has reason to love them qua person, but one has reason not to love them qua friend.

love an entity that is not a subject of well-being, one would be desiring the promotion of well-being of an entity that is not a subject of well-being. The object of one's desire then would be something impossible, like a desire for the number two to be an odd number. Such a desire, however, is incoherent and therefore a desire that one should not have. A minimal requirement for an entity to be an object of love, then, is that the entity be a subject of well-being. The scope of *universal* love would extend to all such beings. Further, all beings that are subjects of well-being plausibly have moral standing. On my view, then, the reasons that ground universal love are the reasons that ground moral standing.[7] The fact that an entity has moral standing gives one reason to include it in the scope of universal love. Similarly, the fact that an entity lacks moral standing gives one reason not to include it in the scope of universal love.

A being with moral standing is an entity that matters morally for its own sake because it is a subject of well-being. For instance, persons and dogs have moral standing because they are subjects of well-being, while rocks and pens do not have moral standing because they are not subjects of well-being. I assume, following Shelly Kagan, that sentience and agency are both individually sufficient, and jointly necessary, grounds for moral standing (Kagan 2019, Ch.1). Here "sentience" means the capacity to feel pleasure and pain, while "agency" means the capacity to have preferences and to act upon those preference. Kagan's view, then, is that an entity *x* has moral standing if and only if *x* has the capacity for sentience or agency. It follows that persons and sentient animals have moral standing, but rocks and pens lack moral standing. I take Kagan's account of moral

[7] To be clear, my substantive claim is that there are normative reasons for universal love. The claim that these reasons are the reasons that ground moral standing is a stipulative claim, not a substantive one. One could have other stipulative conceptions of universal love that appeal to different reasons.

standing as a working assumption. The precise account of the grounds of moral standing is not central to the moral theory sketched in this dissertation.

3.1.2 The Scope of Universal Love

Philosophers writing on love typically assume that love is by nature *particular*. For instance, Michael Slote says: "When we love someone, our concern for them remains separate from other concerns and in that measure is more *particular(izing)* (or in some sense more *personal*) than sheer humanitarianism…" (Slote 2003:116-117). Similarly, Harry Frankfurt claims that "…with regard to love, indifference to the identity of the object of concern is out of question." (Frankfurt, 1999) The idea is that one can only love *particular* individuals with whom one is acquainted, because love treats the beloved as irreplaceable. This rules out, among other things, the possibility of loving an individual qua member of a group. For, if one were to love someone qua member of a group, presumably they could be replaced by any other member of the group as the object of one's love. This leads Frankfurt and Slote to make a sharp distinction between universal love and the humanitarian concerns of universal benevolence.

Now, if love requires acquaintance with a particular individual, universal love will turn out to be rather limited in scope. For any human moral agent, one is not acquainted with a vast majority of beings with moral standing. Given our cognitive limitations, we are only acquainted with a small subset of beings with moral standing. However, I do not think that the object of love must be a particular individual. One can arguably love someone qua member of a group. For instance, one might plausibly love someone qua a compatriot. One might similarly love someone because of some professional role they play — say because they are an artist, an academic, or a priest. My

claim is that it is *possible* to love some individual because of their membership in some group.

Here, I propose a distinction between particular and general love.[8] Particular love is love for some particular individual with whom one is acquainted. General love, on the other hand, is love for an individual who is a member of some relevant group. The difference between particular and general love is that the object of particular love is some particular individual, while the object of general love is *any* individual who is a member of the relevant group. In other words, particular love is *de re* attitude while general love is a *de dicto* attitude. For instance, Bob's particular love for Alice is a case of *de re* love. On the other, Bob's general love for Alice qua compatriot is a case of *de dicto* love. For, if Bob learns that Alice is not a compatriot after all, Bob will not have *de dicto* love for Alice qua compatriot. However, the new knowledge would make no difference to Bob's *de re* love for Alice.

Why think that general love is possible? I want to suggest two reasons in support of this claim. The first is that general love is possible given my account of love. The second is that general love is possible because general hate is possible.

Consider first the possibility of general love given my account of love. On my account, to love someone is to desire the promotion of their well-being for their own sake. Given this account, general love for members of a group would be the desire to promote the well-being of any member of the group for their own sake. For instance, general love for one's compatriots would be the desire to promote the well-being of any

[8] See White (2025) for a similar distinction.

compatriot for their own sake. Importantly, such a general love for one's compatriots is distinct from the love of one's country as such. For, it is possible that one could love one's country and desire to promote national interests even if this imposes significant costs on all of one's compatriots. General love for one's compatriots, on the other hand, is the desire to promote any compatriot's well-being for their own sake.

Consider next phenomena like racial or xenophobic hate. Such phenomena give us reason to believe that general love is possible, independent of my account of love. A plausible way to understand racial or xenophobic hate is that they are attitudes of hatred toward any individual who is a member of a racial or national group. Unfortunately, such hatred often manifests itself in verbal or physical attacks on individuals of the relevant group. When an individual is a victim of such a hate-crime, they are attacked not because of anything about their individual characteristics, but rather, because of their membership in the relevant group. Such hate-crimes are plausibly understood as expressions of hate towards any member of the relevant group. It need not be the case that one hates the individual qua individual, rather one might hate the individual qua member of the relevant group. We might make a distinction then between general and particular hate. Further, if general hate is plausible, then general love is possible as well. For, if one can have an attitude of hatred towards *x*, plausibly one can also have an attitude of love towards *x*.

Nonetheless, one might object, following Frankfurt, that what I am calling general love is not genuine love because general love is indifferent to the *identity* of the object. For instance, assume that Bob is one's compatriot. General love for one's compatriots might manifest itself as love for Bob. But it might as well manifest itself as love for any

other compatriot. The objection goes that so-called general love is not love after all because the object of love is not so fungible. If one truly has an attitude of *love* for Bob, the object of one's attitude could not be replaced by someone other than Bob. Or so the objection goes.

In response, I want to argue that general love can capture the element of irreplaceability that is pre-theoretically characteristic of love. While general love for members of a group does not require love for any particular member of the group, it does not follow that members of the group are therefore replaceable. To think that general love treats the members of the group as replaceable confuses general love for members of a group with love for the group as such qua an abstract object. To see this point, consider the following scenario. Bob is a fellow combatant on the battlefield. He has been injured and needs to be rescued. Is general love for one's compatriots consistent with treating Bob as replaceable and neglecting his need for rescue? Arguably not. If one were to treat Bob as replaceable and not have a *pro tanto* disposition to rescue him, one does not have general love for him. For, if one does have general love for one's compatriots, one would have the *pro tanto* disposition to rescue Bob because he is one's compatriot. To be clear, if Alice were one's fellow compatriot on the battlefield, rather than Bob, one would have the same *de dicto* attitude of love for one's compatriot concerning Alice, rather than Bob. But whether it is Alice or Bob, general love for either is incompatible with treating them as replaceable or fungible with any other compatriot. In other words, while *ex ante* one might have loved Alice, rather than Bob, one cannot *ex post* simply replace Bob with Alice as the object of one's love.

The important upshot of the possibility of general love is that is expands the scope of beings one can love to those with which one is not acquainted. For, if general love is love for someone qua member of a group, all that is necessary for general love for an *F* is a veridical belief of the existence of the group of *F*s. Indeed, one need not be acquainted with any members of the group at all. For instance, one can love factory-farmed chickens without ever having been acquainted with any factory-farmed chicken. Similarly, one could love all beings with moral standing without being acquainted with most of them. All one needs to have knowledge of is the existence of the relevant group. And if Kagan is right, then the relevant group is beings with the capacity for sentience or agency.

Nonetheless, given the cognitive limitations of human moral agents, there may be groups of beings with moral standing of whose existence we are not aware. For instance, there may be rational beings on distant planets whose existence we are unaware. Such beings are outside the scope of universal love because of our lack of cognitive access to them. I suggest, then, that the scope of universal love is all beings with moral standing that one veridically believes to exist. This includes both beings that are objects of one's belief *de re* or as well as those that are objects of one's belief *de dicto*.

Before moving on I want to address the worry raised by Frankfurt and Slote that if love is not particular, there would be no distinction between universal love and the humanitarian concerns of universal benevolence. If universal benevolence is understood as concern for promotion of everyone's well-being for their own sake, then I am happy to say that universal benevolence just *is* universal love. This should not be surprising since on my view love *is* the desire to promote well-being of the beloved. Of course, there are additional behavioral, affective, and cognitive dispositions associated with different

species of love. For instance, there are certain affections and behaviors associated with romantic love and rather different affections and behaviors associated with parental love. What is common to romantic, parental love or any other kind of love, however, is the benevolent desire to promote the well-being of the beloved for their own sake. It should not be surprising, then, if universal love just is universal benevolence.

3.1.3 The Impartiality of Universal Love

On my account, the scope of universal love includes those one in a close relationship with, like one's family or friends, but is not restricted to them. Universal love goes beyond a partial concern for one's close ones and includes strangers in its scope. The impartiality of universal love consists in the recognition that one has reason to love those beyond one's own narrow circle of concerns. The impartiality of universal love, then, consists both in the agent's disposition to love everyone with moral standing and also in the fact that everyone with moral standing is lovable.

I should clarify two points here. First, when I say that everyone is lovable, I do not mean that say that everyone is lovable in the same way. I am not committed, for instance, to the claim that everyone is lovable qua friend or romantic partner. That would be a highly implausible position, because for any given agent there are bound to be many people she would not want to be friends with or romantic partners with. What I committed to is the claim that everyone is lovable, in one way or another. As such, the normative reasons for particular love are plausibly more demanding than the reasons for general love. So, for the vast majority of individuals that one does not have reason to love them qua friend or romantic partner, one would instead have reason for general love toward them, qua person or human being. The claim that everyone is lovable is the claim

that there is *some* reason to love all beings with moral standing (that one believes to exist), not that one has *identical* reasons to love all such beings.

Second, impartiality does not entail equality. I understand the impartiality of universal love to consist in the expansion of the scope of love beyond one's narrow circle of concern. But this does not require that one care about strangers as much as, say, one's children or friends. Plausibly, the desire constitutive of particular loves are stronger than the desires constitutive of general loves. Importantly, this is plausibly justified because the normative reasons for particular loves are arguably stronger than those of general loves. For instance, the fact that an agent has reason to love her children as well as children in Bangladesh does not entail that she has reason to love them all equally. Plausibly, she is justified in loving her children more than children in Bangladesh because of the special relationship she has with her children. Indeed, it would arguably be *wrong* for her to have equal love for her children and children in Bangladesh.

Universal love, then, does not require that one give *equal* consideration to the well-being of all beings with moral standing (that one veridically believes to exist).[9] The impartiality of universal love is compatible with having a stronger desire for the promotion of well-being of one's close ones than the desire for the promotion of the well-being of strangers. Universal love is the desire for the promotion of the well-being of all

[9] I am setting aside the issue of a hierarchy of moral status. If among the beings with moral standing some count morally more than others, the impartiality of universal love will arguably require differential concern and therefore unequal treatment. My focus is on the claim that even within the class of beings with equal moral status, the impartiality of universal love is compatible with differential concern and therefore unequal treatment.

beings with moral standing (that one veridically believes to exist), but this desire may be stronger for some beings than others.

3.1.4 The Unconditional Nature of Universal Love

The last issue I want to address is the *unconditional* nature of universal love. I want to start by clarifying that universal love, on my account, is not unconditional in the strictest sense of the term. For there are normative reasons for universal love on my view. This entails that to be included in the scope of universal love an entity must satisfy certain conditions. In particular, the condition is that an entity must have moral standing. It follows that if an entity lacks moral standing, it would not be included within the scope of universal love. Universal love, then, is conditional on having moral standing.

Nonetheless, universal love is unconditional in a weaker sense of the term. Universal love is unconditional insofar as it does not require a being with moral standing to have any particularly admirable qualities to be included in the scope of universal love. Rather, having moral standing is the only condition one has to satisfy — no being with moral standing can be justifiably excluded from the scope of universal love. To be clear, echoing the point made above, one may have reason to not love a being with moral standing in some particular way, say qua romantic partner, but one would nonetheless have reason to love them, say qua person. Universal love is compatible with change in the *kind* of love one has for another, but it is not compatible with ceasing to love them altogether.

The unconditionality of universal love entails that a person who has done even something seriously wrong still deserves love, by virtue of their moral standing. The

wrongdoing may give one reason to not be friends with the person, but it does not give one reason to not love them at all. The wrongdoing may also give one reason to weaken one's love for them. Nonetheless, the fact that the person has moral standing gives one reason to love them, qua person, despite their wrongdoing.

3.2 Christine Swanton on Universal Love

Let me close this chapter by discussing Christine Swanton's account of universal love (Swanton 2003). I want to briefly discuss her account in order to further clarify my own account. I focus on her account in particular because it is one of the most developed accounts of universal love in the literature.

On Swanton's account, universal love is the *expression* of an agent's loving nature which manifests itself in the disposition to form bonds with anyone[10]. Briefly, on her account universal love has the following features: it is (i) a form of coming close, i.e., it seeks to form bonds with the objects of universal love, (ii) particular, i.e., the objects of universal love are particular individuals, (iii) universal, i.e., the objects of universal love could be anyone, (iv) impartial, i.e., universal love transcends partial concerns for one's family and friends, (v) unconditional, i.e., there are no 'acceptance' and 'termination' conditions for universal love (Swanton 2003:117-127). The salient differences between our accounts are that I reject features (i)-(iii).

In the remainder of this section, I first briefly explain each of the features and then highlight the salient differences between our accounts of universal love.

[10] Or anything. Unlike my usage of "love" which is restricted to "love for someone," Swanton uses "love" in a wider sense that have as its object not just people or animals but also things like mountains.

3.2.1 Swanton's Account

On Swanton's view, universal love is a form of 'coming close'. What does it mean to 'come close' to another? She fleshes it out in terms of bonds. The idea is that to love someone is to have a bond with them, for instance a bond of friendship or romance. Such bonds presumably include both a history of interaction as well as strong emotional attachments and vulnerabilities. On Swanton's view, the existence of bonds distinguishes love from other attitudes like respect which do not involve a bond of coming close.

To love someone, then, requires a bond with them, be it a bond of romance or friendship, or a familial bond. It is possible to form a bond, however, only with *particular* individuals. On Swanton's view, the object of love is always a particular individual. Importantly, on her view the particularity of love distinguishes universal love from the humanitarian concerns of universal benevolence. The key difference, on Swanton's view, is that universal benevolence does not require forming a bond with any particular individuals while universal love does.

Further, Swanton suggests that universal love does not require any property or properties that ground universal love and thereby determine its scope. In other words, there are no *reasons* for universal love that determine its rational scope. Rather, universal love *expresses* an agent's loving nature. The *universality* of universal love consists in the fact that the agent is open to loving everyone and does not limit her love to her some class of beings.

Universal love, then, is also impartial on Swanton's view insofar as it is not restricted to a partial concern restricted to, say, one's family or friends but extends to

people and animals that are strangers to oneself. As with universality, the impartial nature of universal love has to do with the loving nature of the agent rather than the lovability of the beloved. On her view, what is impartial about universal love is the fact that the agent is open to forming bonds beyond her own narrow circle of concern.

Lastly, on Swanton's view universal love is unconditional. By the unconditionality of universal love, she means that it "has neither 'acceptance' nor 'terminating' conditions" (Swanton 2003:125). In other words, there are no conditions that someone has to meet in order to be an object of universal love, nor are there any conditions which would lead an agent who manifests universal love to cease loving them.

3.2.2 Contra Swanton's Account

My account of universal love differs from Swanton's account in three important respects. First, unlike Swanton's expressive account, on my account there are normative reasons for universal love which determine its scope. In particular, I suggest that the scope of universal love is all beings with moral standing (that one veridically believes to exist). On my view, the reasons that justify universal love just are the reasons that ground moral standing. Contra Swanton, universal love on my account is not merely an expression of an agent's loving nature but is an agent's loving disposition as a fitting response to the fact that she inhabits a world with beings with moral standing. In particular, the fact that someone has the capacity for sentience or agency gives one reason to love them.

The difference between Swanton's expressive view and my reasons-view entails a difference in how we understand the unconditionality of universal love. On Swanton's account, universal love is unconditional in the strict sense. For, on her account universal

love expresses an agent's loving nature and is not as such constrained by any facts about the object of the agent's love. By contrast, on my account, the existence of reasons for universal love entails that universal love is not strictly unconditional. For there is a condition an entity must meet in order to be an object of universal love — it must have moral standing. Nonetheless, universal love is an unconditional from of love, in a weaker sense, since there are no further conditions an entity with moral standing must meet in order to be within the scope of universal love. Importantly, there are no conditions under which an entity with moral standing deserves *not* to be loved.

Second, unlike Swanton's account, on my account coming close is not *essential* to universal love. I agree with Swanton that love paradigmatically manifests itself in the bonds of friendship, marriage, and the like. But arguably such coming is not essential to love. What is essential to love is the desire for the promotion of the well-being of the beloved for their own sake. The promotion of someone's well-being however does not require forming a bond with them. Indeed, in some cases coming close to the beloved could reduce their well-being.

Consider the following case. Imagine that Eve is a new mother whose marriage has just fallen apart, and she now finds herself in a dire economic and emotional state for the foreseeable future. She reasons that it would be in her infant's best interests to be raised by a more stable family. Accordingly, she decides to put her infant up for adoption, despite much grief to herself. If Swanton is right Eve would presumably not be acting out of love. For Eve intends to do quite the opposite of coming close, she wants, *not* to bond with her infant, but instead to allow the infant to share a bond with another mother. Contra Swanson, however, Eve would be acting out of love. My account can explain that

Eve would be acting out of love for her infant because she would be acting on her desire to promote her infant's well-being even though it causes her much grief to do so.

Third, unlike Swanton's particularized account, on my account universal love is not restricted to love of particular individuals but includes general love for beings with moral standing with whom one is not personally acquainted. This point is related to the second point above. If universal love does not require forming personal bonds, then there is no necessity that the object of universal love be *particular* individuals. Rather the object of universal love can be individuals that one veridically believes to exist by virtue of their membership in some relevant group. On my account, the relevant group is the class of beings with moral standing. Further, one can love beings in this class by desiring the promotion of their well-being for their own sake.

The difference between Swanton's particularized view and my more permissive view entails a difference in how universal love can be manifested. On Swanton's view, an agent who manifests universal love would form bonds with the beings with which she is acquainted. However, any human moral agent is only acquainted with a limited number of beings so the number of beings the agent would in fact love would be rather limited. On my view, on the other hand, an agent who manifests universal love would desire the well-being of all beings with moral standing that she veridically believes to exist. Importantly, this is not limited to beings she is acquainted with but also those she is aware of only through description.

To sum up the chapter: Universal love is the desire for the promotion of the well-being of all beings with moral standing. Importantly, this includes beings one is not acquainted with and does not require equal love for all.

4 Universal Love and Moral Theory

I have so far developed an account of universal love as the desire for the promotion of the well-being of all beings with moral standing (that one veridically believes to exist). I now turn to the central issue of this dissertation: What role can universal love play in moral theory? I discuss this issue in this chapter and the next.

In this chapter, I first argue that universal love cannot be the object of a moral requirement because there cannot be a *duty* to love. Briefly, there cannot be a duty to love because love is not within an agent's voluntary control. I then suggest that the disposition to love universally is a moral virtue. Briefly, the disposition to love all beings with moral standing is a morally excellent disposition. The upshot of these claims is that universal love can play a foundational role in a virtue ethical moral theory, but not in a deontological moral theory.

In the next chapter, I develop a virtue ethical moral theory in which the moral virtue of universal love plays the central role of *unifying* the (other) moral virtues. Consequently, universal love also plays a central role in my account of moral reasons for action.

4.1 Preliminaries

Our present interest is to clarify the role that universal love can play in moral theory. In this section, I lay the groundwork for what follows in the subsequent sections of this chapter. I start by clarifying my basic assumptions about morality as a normative domain. I then provide a brief overview of the major approaches in moral theory and the role that universal love could play on each of these approaches.

4.1.1 Normative Domains and Reasons

I take morality to be a normative domain, among others, which allows us to evaluate actions, beliefs, or desires. Following Dale Dorsey (2016:9), we may understand normative domain in *functional* terms: it takes as inputs certain facts about an action, belief or desire and generates a corresponding normative output of "good," "bad," "required," "permissible," etc.

Any given normative domain, then, generates *domain-relative reasons* to perform an action, hold a belief or have a desire. For instance, the fact that one might harm an innocent person if one were to drive drunk gives one a *moral* reason not to drive while drunk. Similarly, the fact that one might harm oneself or damage one's car gives one a *prudential* reason not to drive while drunk. These domain-relative reasons may be deontic or evaluative. Relative a given domain, deontic reasons specify what beliefs, desires or actions are permissible or required, while evaluative reasons specify the *value* of beliefs, desires, or actions. An important difference between deontic and evaluative reasons is that deontic concepts do not admit of degrees while evaluative concepts do. For instance, deontic claims about an action may specify whether or not the action is permissible. On the other hand, evaluative claims about an action may specify the degree of desirability (goodness) or undesirability (badness) of the action.

A *practical* normative domain takes as inputs certain facts about an action (or a disposition to act) and generates evaluative (e.g., good, bad) and deontic (e.g., required, permissible) outputs about the action (or the disposition to act). Leaving open the possibility of other practical normative domains, I assume that morality and prudence are at least two distinct practical normative domains. In particular, I take prudence to be a

self-regarding domain concerned with the agent's own interests and morality to be an other-regarding domain concerned with the interests of others. This is not to say that there is no overlap between morality and prudence. Often, promoting others' interests is also in one's own interest. However, morality and prudence are distinct domains that sometimes conflict, even if in many cases (perhaps even the majority) they are in harmony.

Two clarifications are in order here. First, I am not suggesting that practical normative domains like morality and prudence cannot be used to assess beliefs or desires. Plausibly, racist and xenophobic beliefs — roughly, beliefs that members of some racial or national group are inferior to members of one's own racial or national group — are morally bad. Arguably, such beliefs are not merely instrumentally bad but also intrinsically bad because they devalue some members of a group based on morally irrelevant features. Similarly, the belief that smoking two packs of cigarette every day is perfectly healthy is prudentially bad.[11] So, practical normative domains like morality and prudence assess actions but they may also assess beliefs and desires.

Secondly, in the discussion that follows I have human agents in mind. There may be rational agents that are radically different from human agents and many of my claims below would not apply to them. When discussing the possibility of a duty to love, for instance, the main argument I discuss appeals to human psychological limitations. However, there may be rational agents who do not share human psychological

[11] In this case it is also epistemically bad insofar as it is not adequately sensitive to the evidence concerning the ill-effects of smoking.

limitations, and these considerations will not apply to them.[12] Similarly, the *kind* of agent one happens to be might make a difference to what the relevant virtues are for the agent. I will focus on human agents in what follows, setting aside the issue of what one might want to say about non-human agents.

4.1.2 Universal Love and Moral Theory

I take a moral theory to be a theory of the *grounds* of moral properties — the moral properties of actions, intentions, character traits, practices, institutions, etc.[13] That is, a moral theory specifies the properties which determine moral properties like moral goodness and moral permissibility. Further, by specifying the grounds of moral properties, a moral theory also explains *why* the moral domain generates the normative outputs that it does in response to certain factual inputs. For instance, divine command theories say that moral properties are grounded in conformity to God's will. On such theories, conformity to God's will determines and thus explains moral properties.

Of particular interest, given the *practical* nature of morality, is a moral theory's account of moral reasons for action — the properties which determine, and thereby explain, the moral permissibility and moral goodness of actions. For instance, divine command theories say that an action is morally impermissible, for instance, if it is forbidden by God. In general, a central task of any plausible moral theory is to provide an

[12] Consequently, such beings may have a duty to love because they can love someone at will. At least, the commandability objection I discuss in the next section will not apply for such beings.
[13] This is, of course, a conception of an exhaustive moral theory. There might be more theories more limited in scope — say only a theory of the deontic moral properties of actions.

account of moral reasons for action, since a crucial concern of morality is to help answer the question of how one should act.[14]

Now, our interest in this chapter is to investigate the role that universal love can play in a moral theory. To that end, it will be useful to briefly consider what role universal love can play on the dominant approaches to moral theory — deontology, virtue ethics and consequentialism. Roughly, while deontological theories ground moral properties in basic moral duties and virtue ethical theories in the moral virtues, consequentialist theories ground moral properties in the moral goodness of states of affairs. Let us briefly consider each approach in turn.

First, on a deontological approach, there might be a *fundamental* moral duty to love universally from which other (non-fundamental) moral duties are derived. On this approach, the ground of morality would be that all moral agents are obligated to love all beings with moral standing (that one veridically believes to exist). The duties not to murder or torture, for instance, would be determined and explained by this fundamental moral duty. A strong version of this deontological approach would say that loving universally love is the only fundamental moral duty. An alternative pluralist approach would leave room for other fundamental moral duties.

Second, on a virtue ethical approach, the disposition to love universally might be a moral virtue. There are two ways of spelling this out. A modest view would be that the disposition to love universally is *one* of the foundations of morality, along with all the

[14] To be clear, this leaves open the nature of moral reasons for action. For instance, one might be skeptical of moral obligations and only want to speak in terms of the moral goodness or badness of actions. Similarly, one might be skeptical about the prospects of reducing moral reasons for action to one or more general moral principles.

other moral virtues like generosity, benevolence, justice, and fairness. On this approach, universal love is *pro tanto* relevant to determining moral properties. Call this the weak virtue ethical approach. A more ambitious view would be that the virtue of universal love is foundational to morality in a stronger sense. On this latter view, the disposition to love universally is not just one moral virtue among others that is *pro tanto* relevant to determining moral properties. It is, rather, the sole foundational moral virtue that is *conclusive* for determining moral properties. Call this the strong virtue ethical approach. The moral theory I develop in the next chapter takes this strong virtue ethical approach. As I explain in the next chapter, universal love is the sole foundational moral virtue insofar as it *unifies* the moral virtues.

Third, on a consequentialist approach, there are two main possibilities. First, universal love might be instrumentally relevant for morality insofar as cultivating universal love helps moral agents to reliably promote the good. On this approach what matters is the production of good states of affairs and universal love is instrumentally valuable to that end. Second, universal love might be an intrinsically valuable second-order good. On this view, a character trait like the disposition to love universally might be a second-order intrinsic good because it is an intrinsically good disposition directed towards first-order intrinsic goods — the well-being of beings with moral standing.

Having briefly introduced the two possible consequentialist approaches, I want to quickly set them aside because neither of the approaches is *prima facie* promising. On the first approach, universal love is only instrumentally good from the moral point of view. It is the goodness of states of affairs that grounds moral properties on this approach. On the second approach, universal love would be intrinsically good from the moral point of view

insofar as it grounds the moral properties of at least some character traits. However, even on the second approach, universal love would play no role in explaining moral reasons for action. This is because, like the first approach, the moral properties of actions would presumably be determined by the causal relationship between actions and the first-order goods (i.e., well-being). Neither of these consequentialist approaches, then, is promising for giving universal love a foundational role.

Accordingly, in the remainder of this chapter I set aside the consequentialist approach and focus instead on the deontological and the virtue ethical approaches. Ultimately, I reject the deontological approach in favor of the virtue ethical approach.

4.2 Love and Duty

Consider first a deontological moral theory on which universal love is the object of a fundamental moral duty. The prospects of such a theory depend crucially on the issue of whether there can be a *duty* to love in the first place. The issue here is whether it is conceptually possible for there to be a normative *requirement* to love someone. For, if it is not possible, then universal love cannot be the object of a moral duty, fundamental or otherwise. Roughly, the objection is that love is not within an agent's voluntary control, and there can be no duty to love in the absence of such voluntary control. Call this the *commandability objection*.

In this section, I discuss the commandability objection. I first spell out the objection in detail. I then consider a recent response to it and argue that the response is inadequate. I conclude that a duty to love is conceptually impossible. The upshot of this conclusion is that love of any kind, including universal love, cannot play a foundational role in a deontological moral theory.

4.2.1 The Commandability Objection

The commandability objection to the possibility of a duty to love is that love is not within an agent's voluntary control such that she can bring it about at will. Here is how we can flesh out the objection. An agent *A* has a duty to *X* only if *A* has voluntary control over *X*. Call this the *commandability principle*. Why couldn't there be duties that are not within an agent's rational control? One plausible answer, following Stephen Darwall (2009), is that duties that are not fulfilled warrant blame. But an agent could not be justifiably blamed for something outside of her voluntary control. It follows that there could not be duties that are not within an agent's voluntary control, at least if there is a conceptual connection between duties and blameworthiness. Even if one does not endorse this rationale, the general principle of *ought implies can* is intuitive enough to take the commandability objection seriously.

We must distinguish, first of all, between *direct* and *indirect* voluntary control. Let us say that an agent *A* has direct (full) control over *X* at time t if and only if (1) *A* can make a volition now that reliably brings about *X*, and (2) *A* can make a volition now that reliably brings about not-*X*. For instance, I plausibly have direct voluntary control now over the act of scratching my forehead right now — I could either do so or refrain from doing so. On the other hand, there may be things over which I do not have such direct voluntary control but have *indirect* voluntary control. For instance, I plausibly have indirect voluntary control over the state of affairs that I lose five pounds. I do not have direct voluntary control in this case because I could not bring it about right now that I lose five pounds now. But plausibly I do have indirect voluntary control over losing five pounds in the near future — I could start exercising, get on a diet, etc. that will eventually

bring it about that I lose five pounds. In general, an agent *A* has indirect (full) control over *X* at time t if and only if (1) *A* does not have direct control over *X* at t, and (2) *A* can make series of volitions (over time) that reliably brings about *X*, and (3) *A* can make a series of volitions (over time) that reliably brings about not-*X*.

A quick clarification about the notion of reliability I have just invoked. Reliably bringing about a state of affairs requires that the probability of one's volition bringing about the state of affairs is high enough, but it does not require the certainty of bringing it about. That is, reliability requires a probability much greater than 0.5, but it does not require a probability of 1. Beyond this, I leave the notion of reliability vague and do not stipulate a precise probability threshold.

We can now state the commandability principle as follows: For any normative domain, an agent *A* has a duty (at time t) to *X* only if *A* has (direct or indirect) voluntary control (at time t) over *X*. That is, duty requires voluntary control, whether direct or indirect.

How does the commandability principle apply in the case of love? The principle would say that one can have a duty to love someone only if one has some form of voluntary control over whether one loves them. Consider now whether there can be a duty to love, given my account of love. Recall that, on my view, to love someone is to desire the promotion of their well-being for their own sake. Further, on my view, to desire *X* is to be disposed to exercise one's agency to realize *X*. Arguably, given my account of love, there could *not* be a duty to love, because one's desires are not within one's voluntary control.

To begin with, one's desires are not within one's direct voluntary control. Consider, for instance, a mundane desire like the desire to eat pizza for lunch. One could not desire to eat pizza at will, for one could not bring it about at will whether or not one desires to eat pizza for lunch. Rather, one either has or lacks a desire to eat pizza that one discovers through introspection. Similarly, the desire constitutive of love is not within one's directly voluntary control, rather one discovers whether or not one loves someone through introspection. To be clear, given my account of love, what such introspection would reveal is not a *feeling* but a disposition of the will to want the desired object.

What about indirect voluntary control over one's desires? Are there any practices that could allow one to cultivate a desire over time? Plausibly, one way to cultivate a desire is to start acting *as if* one did have the desire. For instance, imagine that Sarah wants to cultivate the desire to eat vegetables. Suppose that Sarah is currently rather averse to eating vegetables but knows that eating vegetables is good for her health. She wants to cultivate the desire to eat vegetables in order to improve her health. She cannot bring it about that she desires eating vegetables right away. However, she could try to cultivate the desire by starting to eat vegetables. As she continues to do so, she might eventually acquire the disposition to eat vegetables as a result of habituation. Nonetheless, it is doubtful that this process of forming a desire through habituation has a high enough probability of success to amount to *control* over the desire.

Habituation plausibly increases the likelihood of forming the desire compared to no habituation, but it is highly doubtful that it raises the probability beyond the threshold required for a *reliable* process of desire formation. As I clarified above, control over X does not require that one have the ability to bring about X with certainty of success (i.e.,

probability=1). Nonetheless, control over *X* does require the ability to bring about *X* in a reliable manner (i.e., probability >> 0.5). Without getting more precise, the relevant reliability is the ability to bring about *X* with a high enough probability of success. But it is highly doubtful that practices like habituation give one such ability with respect to one's desires. At best, such practices might raise the probability of having certain desires, relative to the probability of having these desires in the absence of these practices. However, it is highly doubtful that these practices raise the probability of having these desires above the threshold required for the reliability needed for control over one's desires. So, if love is some sort of desire, there could not be a duty to love.

To sum up: If love is the desire for someone's well-being for their own sake, love is neither within one's direct voluntary control nor is it within one's indirect voluntary control. It follows that there could be no *duty* to love.

4.2.2 Matthew Liao's Response to the Commandability Objection

I have so far developed the argument that that there can be no duty to love because love is not sufficiently within an agent's control. Matthew Liao, however, has recently defended the possibility of a duty to love against the commandability objection. He argues that love is in fact commandable and therefore there can be a duty to love (Liao 2006, 2015). I briefly discuss his argument and then argue that it does not establish the commandability of love.

Liao argues that there are three ways by which an agent can bring it about that she loves someone — internal control, external control, and cultivation of one's general capacity to love (Liao 2015, Ch. 4). In particular, he argues that love is commandable

because one has voluntary control over one's emotions as well as one's valuational attitudes. While his focus on control over one's emotions and valuational attitudes, his arguments can plausibly be extended to one's desires as well. Let us, then, briefly consider the three methods to control one's emotions and valuational attitudes proposed by Liao. In illustrating the methods, I shall paraphrase the examples Liao himself provides.

The method of *internal control* involves giving oneself reasons to have a particular emotion or valuational attitude. For instance, suppose that I have been invited to a close friend's wedding this evening, but I am in a bad mood. Knowing that my friend would want me in a good mood, I could tell myself that I should be joyful on this special occasion. It is likely that I would not be in a bad mood anymore and be joyful instead. Similarly, if one wants to have a more positive attitude towards the environment, one could tell oneself the benefits of a better environment like improved quality of life. It is likely that one would have a more positive attitude towards the environment as a result.

The method of *external control* involves placing oneself in circumstances that elicit a particular emotion or valuational attitude. For instance, if I tend to become angry when there is no food in the refrigerator, one way I can control my anger is by ensuring that my refrigerator is always stocked. Similarly, if one wants to have a more negative attitude towards the lack of charitable aid to people living in extreme poverty, one could place oneself in areas where people live in extreme poverty.

The method of *cultivation* involves cultivating one's emotional and attitudinal capacities by acting as if one had the emotion or attitude. Cultivation makes it more likely that one will have a particular emotion or attitude at a particular time as a result of

habituation without requiring internal or external control. For instance, one can cultivate one's capacity for joy through joyful behavior, like smiling. Through repeated practice of joyful behavior, it is likely that one would cultivate one's capacity for joy. One could similarly cultivate a positive or negative attitude through repeatedly acting as if one had the attitude.

Can these methods be applied to desires? Recall that on my view a desire for *X* is a disposition to act in ways that one believes are likely to bring about *X*. The most promising method to bring it about that one has a desire would be the method of cultivation. Consider again the example considered above of Sarah, who wants to have the desire to eat vegetables. The most promising method to bring it about that she has the desire would be to start eating vegetables. Further, she could use the methods of internal and external control to supplement the method of cultivation. When she finds it difficult to motivate herself to buy vegetables, she could use the methods of internal or external control as psychological crutches. For instance, she could tell herself the benefits of eating vegetables (internal control). Or she could go to a local health food store, knowing that it makes her more likely to eat vegetables (external control). The three methods Liao suggests, then, could certainly be applied to desires.

The pertinent question is whether cultivation, supplemented by internal and external control, is sufficient for an agent to have voluntary control over her desires. Arguably, it is not. As I suggested earlier, cultivation of desires through habituation plausibly increases the likelihood of forming the desire compared to no habituation, but it is highly doubtful that it raises the probability beyond the threshold required for a *reliable* process of desire formation. Even when supplemented by internal and external control,

the process remains unreliable because the likelihood of success in bringing about the intended desire is not sufficiently high. This means that there may well be cases where one successfully cultivates a desire, nonetheless one did not have control over the desire because the process that produced the desire was not reliable.

To see this point, we must draw a distinction between indirect voluntary *control* and indirect voluntary *influence*. For Liao's argument to be successful it must show that an agent has indirect voluntary control over her desires. Recall that, if at time t an agent *A* has indirect voluntary control over *X*, then she can reliably bring it about whether *X* or not *X* at some future time t'. For instance, an agent has indirect voluntary control over losing five pounds because she can reliably bring it about whether or not she loses five pounds in the future. Of course, the control is limited in terms of how quicky she does so. Within certain limits, however, she does have indirect voluntary control by virtue of her control over her exercise and eating habits. On the other hand, it is doubtful that Sarah has similar control over her desire to eat vegetables. For it is highly probable that despite her best efforts to cultivate the desire over an extended period of time, she nonetheless despises eating vegetables. What Sarah's efforts do is *influence* her desire to eat vegetables — her efforts make it more likely that she has the desire to eat vegetables compared to the likelihood of her having the desire in the absence of her efforts.

Liao's argument, then, only establishes that an agent can influence whether she loves someone, but it does not establish that she can control whether she loves them. Importantly, given the absence of control, an agent cannot be *obligated* to bring it about that one loves someone. At best, one could have a duty to *cultivate* love for someone insofar as doing so increases the likelihood that one loves them.

The upshot of the discussion so far is that there can be no duty to love because (human) agents lack voluntary control over whether they love someone. Consequently, universal love cannot play a foundational role in a deontological moral theory.

4.3 Universal Love as a Moral Virtue

I have so far argued that universal love cannot play a foundational role in a deontological moral theory because love cannot be the object of a normative requirement. I now want to suggest that universal love can play a foundational role in a virtue ethical moral theory because the disposition to love universally is a morally excellent disposition.

I begin by sketching an account of moral virtue based on Christine Swanton's target-centered account of virtue. Using this target-centered account of virtue, I then sketch an account of universal love as a moral virtue.

4.3.1 A Target-Centered Account of Virtue

I take the moral virtues to be a proper subset of the practical virtues — virtues that concern action as opposed to, say, intellectual virtues that concern belief. We can distinguish *kinds* of practical virtues corresponding to the different practical normative domains. As discussed in section 4.1 above, morality and prudence are at least two distinct practical normative domains. Accordingly, the practical virtues include at least the moral and prudential virtues. Broadly speaking, what distinguishes moral virtues from prudential virtues is that the former are other-regarding virtues while the latter are self-regarding virtues. As with the general relationship between morality and prudence, there is often harmony between the moral and prudential virtues though there can be occasional conflicts between them.

I shall now sketch an account of moral virtue within a broader account of practical virtue that builds on Christine Swanton's target-centered account of virtue (Swanton 2003). Central to her target-centered account of virtue are the notions of the *field* and the *target* of a virtue.

The *field* of a virtue consists of the elements that are the sphere of concern of the virtue. That is, the field of a virtue consists of the aspects of the world, broadly construed, that are normatively significantly insofar as an agent should respond to them in a fitting manner in order to possess the virtue. For instance, the field of benevolence is others' well-being, while the field of temperance is one's bodily pleasures. In general, the field of a virtue could include concrete individuals, abstract entities, or state of affairs. The field of a virtue could include concrete individuals like persons, animals, trees, or mountains. It could include abstract entities like knowledge, creativity or beauty. It could also include states of affairs like being a dangerous situation or having maximal well-being.

A *virtue* for a given normative domain is a disposition to respond to elements within its field(s) in an excellent way, relative to that domain. Excellence is to be understood relative to a normative domain. A normative domain gives us a point of view to evaluate dispositions. In particular, moral goodness consists in being conducive, in one way or another, to others' interests or wills. A moral virtue is a morally excellent disposition — a disposition that a maximally good disposition from the moral point of view. The moral virtues are thus constitutive of excellent moral agency.

The *target* of a virtue is what the virtue is aimed at within its field. For instance, the target of benevolence is the promotion of others' well-being. A virtue's field specifies

what the agent should respond to, while the virtue's target specifies *how* the agent should respond. Here, two points must be kept in mind. First, there may be many appropriate ways of responding to an element of a virtue's field — promoting, protecting, respecting etc. Second, the appropriate way to respond is often context dependent. Consider, for instance, the moral virtue of fairness. In general terms we could say that fairness is aimed at treating others fairly. In some cases, treating someone fairly may require respecting their choices, say by honoring someone's choice to not be resuscitated. However, in other cases treating someone fairly may require promoting their interests, say by giving to the needy. In still other cases treating someone fairly may require protecting their interests, say by protecting a vulnerable individual from getting exploited.

This account of moral virtue also gives us an account of moral reasons for action. We can start by defining virtuous action as follows: An act *A* is virtuous with respect to some virtue *V* to the degree that *A* achieves the target of *V*. For instance, an act is virtuous with respect to benevolence as long as it successfully promotes others' well-being in a given circumstance. We can further compare acts to assess their degree of virtuousness with respect to benevolence by assessing how well they promote others' well-being. An act's being virtuous (to whatever degree) gives *pro tanto* domain-relative reason for performing the act (to the degree of its virtuousness). Similarly, an act's being vicious gives *pro tanto* reason against performing the act (proportional to its degree of viciousness). Virtues and vices are thus the grounds of (domain-relative) reasons for actions.[15]

[15] This is no doubt a controversial claim, but I endorse the claim without defending it here.

Two clarifications are in order here. First, it is not necessary for an act to be virtuous that it be produced by a virtuous disposition. For, one may perform an act that is virtuous for the wrong kind of reasons. For instance, one may perform a generous act because of peer pressure or fear of divine judgement, rather than a generous disposition. In general, a virtuous act is a good act, but it may be done for the wrong reasons. Such acts are good because they meet the targets of virtue, without being praiseworthy because they are not done for the right reasons. Second, an act's being produced by a virtuous disposition is not sufficient for the act to be virtuous. For, an act produced by a virtuous disposition may fail to achieve the target of the virtue. For instance, an act produced by a benevolent disposition may not successfully promote others' well-being, perhaps because of bad luck, and fail to achieve the target of benevolence. In general, one may act for the right reasons but fail to perform a good act. The overall point is that the goodness of acts depends on achieving the targets of virtues, which can come apart from the goodness of one's motivations. Consequently, praiseworthiness of acts can come apart from their goodness.

Next, we can generalize from virtuousness with respect to particular virtues to virtuousness with respect to a normative domain: An act *A* is virtuous with respect to some normative domain *D* to the degree that *A* achieves the targets of all the relevant virtues of *D*. It follows that an act is morally virtuous to the degree that it achieves the targets of all the relevant moral virtues. To be clear, I do not assume that the moral virtuousness of an act settles its overall virtuousness. An act *A* is overall virtuous to the degree that *A* achieves the targets of all the relevant virtues of all the relevant normative domains. An act might achieve the targets of the relevant moral virtues but not the

relevant prudential virtues. I leave open the overall normative properties of such acts. My focus, instead, is on the moral properties of acts.

4.3.2 Universal Love as a Moral Virtue

With the target-centered account of virtue in hand, we can now begin to flesh out the proposal that the disposition to love universally is a moral virtue. To begin with, given the account of universal love sketched earlier, the disposition to love universally is the disposition to desire the promotion of the well-being of all beings with moral standing (that one veridically believes to exist) for their own sake. Being an other-oriented disposition concerned with their well-being, the disposition to love universally is clearly a candidate for being a moral virtue.

The pertinent question, then, is whether the disposition to love universally is a morally excellent disposition. Here, I appeal to intuitive moral judgements to make the case that the disposition to love universally is a morally excellent disposition. Consider the following variety of desires one could have concerning the well-being of all beings with moral standing (that one veridically believes to exist). For ease of expression, below I shall use "others" to refer to all beings with moral standing that one veridically believes to exist. First, one could have a desire to harm others. That is, a desire for decrease in others' well-being. This attitude, which we may call universal hatred, is undoubtedly a morally bad or vicious disposition. Second, one could lack any desire whatsoever about others' well-being. Such indifference about others' well-being is plausibly morally neutral. It is not morally bad insofar as one does not desire a decrease in others' well-being but nor is it morally good insofar as one has no concern whatsoever for others' well-being.

Third, one could have a desire to not harm others. That is, one could have a desire to *not decrease* others' well-being. This would involve a disposition to not inflict harm on others. Setting aside the issue of how harms should be aggregated, such an attitude is morally good insofar as it shows at least some positive concern for others' well-being. Fourth, one could have a desire to *protect* others from harm, along with a desire to not harm them. That is, one could desire to (i) not decrease others' well-being, and (ii) *prevent a decrease* in others' well-being. This would involve not just a disposition to refrain from harming others but also a disposition to rescue them threats. This fourth disposition is morally better than the third because it shows a greater positive concern for others' well-being.

Lastly, one could have a desire to *benefit* others, along with the desires to protect them and not harm them. That is, one could desire to (i) not decrease others' well-being, (ii) prevent a decrease in others' well-being, and (iii) increase others' well-being. This would involve not just the dispositions to rescue others from threats and to refrain from inflicting harm on them, but also a disposition to promote their well-being. This last set of dispositions is plausibly the morally best disposition because it shows the greatest positive concern for others' well-being, relative to all the other dispositions considered above. This last set of disposition just is the disposition to love universally. I shall call this the moral virtue of universal love.

To clarify the role that the moral virtue of universal love can play in moral theory, we need a more worked out account of the virtue. On the target-centered conception of virtues, this means that we need to specify the *target* of universal love. To begin with, the field of universal love is the well-being of all beings with moral standing who stand to be

affected, in one way or another, by one's agency. What does universal love say about *how* a moral agent ought to respond to the well-being of all individuals whose well-being is at stake? I suggest that the target of universal love is the promotion of the sum of the *weighted*, *relevant* interests of all beings with moral standing that stand to be affected by one's agency. Let me explain.

Let us say that each individual whose well-being stands to be affected by a moral agent's agency has an interest in being aided and an interest in not being harmed. An interest in not being harmed is an interest in not having one's well-being decreased by the agent. An interest in being aided is an interest in having one's well-being protected (from decrease) or promoted by the agent. On the one hand, universal love does not justify a thoroughgoing aggregation of *all* harms and benefits, since a thoroughgoing aggregation fails to adequately manifest love for *each* individual. A thoroughgoing aggregation treats each individual and their well-being as fungible, i.e., replaceable, rather than treating them as non-fungible, i.e., non-replaceable, individuals whose well-being matter for their own sake. On the other hand, universal love does not justify a thoroughgoing anti-aggregative stance of *never* aggregating any harms and benefits. A thoroughgoing anti-aggregative stance fails to adequately manifest love for *all* individuals. If one truly cares about the well-being of all individuals, one will be open to aggregation in at least some cases. Neither thoroughgoing aggregation nor thoroughgoing anti-aggregation is consistent with a desire for the promotion of well-being of *all* individuals for *their own sake*.

Universal love, I suggest, justifies *partial* aggregation of harms and benefits. That is, universal love justifies aggregation of *relevant* interests (I talk more about the notion

of relevance below). This distinguishes it from the thoroughgoing anti-aggregative view insofar as it is open to aggregation in some cases. However, universal love does *not* justify aggregation of *irrelevant* interests. This distinguishes it from the full-aggregative view insofar as it is *not* open to aggregation in some cases. Only aggregating relevant interests arguably manifests a concern for the promotion of the well-being of all individuals for their own sake.

What makes an interest relevant? I do not take any specific position on the relevance factors that determine relevance but only mention some plausible relevance factors:

1. the moral status of each individual,
2. the current level of well-being of each individuals,
3. the change in the level of well-being of each individual (if aided or harmed),
4. doing vs. allowing harm,
5. whether the agent has a special relationship with any of the individuals involved.

The relevance factors, whatever they are, collectively determine both (i) which interests are *relevant,* and (ii) the *weight* of the relevant interests, i.e., how much each relevant interest matters. The target of universal love is the promotion of the sum of the weighted, relevant interests of all beings with moral standing that stand to be affected by one's agency. Of course, I have so far only provided a very rough sketch of an account of relevance. A fully worked out view would tell us (i) what the relevant factors are, (ii) how much each factor matters, and (iii) provide justification for (i) and (ii). At this stage,

however, I am only providing a rough sketch, leaving open how the details are filled in for future work.

Let us take stock. The upshot of the discussion so far is that universal love can play a foundational role in a virtue ethical moral theory. Recall the distinction I made earlier between the weak and strong virtue ethical approaches depending on whether the disposition to love universally is one moral virtue among others or the sole foundational moral virtue. What I have said so far is compatible with either approach. In the next chapter I develop a strong virtue ethical approach on which the disposition to love universally is the sole foundational moral virtue.

5 The Primacy of Universal Love

In the previous chapter, I argued that universal love can play a foundational role in a virtue ethical moral theory. In this chapter I sketch a moral theory in which the disposition to love universally is the sole cardinal moral virtue that grounds all the other moral virtues. Further, the moral virtue of universal love also *unifies* all the other moral virtues and is thus conclusively relevant for determining moral reasons for action.

I begin by introducing the notion of a cardinal virtue. Briefly, a cardinal virtue is a virtue that grounds some virtues in its normative domain but is not itself grounded by any virtue in its normative domain. I then argue that the disposition to love universally is the sole cardinal moral virtue that grounds all the other moral virtues. I conclude by sketching an account of moral reasons for action. Roughly, the disposition to love universally is conclusively relevant for determining moral reasons for action because it helps adjudicate conflicts between the non-cardinal moral virtues that ground *pro tanto* moral reasons for action.

5.1 Cardinal Virtue and the Unity of the Virtues

The notion of a cardinal virtue is relevant for the important issue of the unity of the virtues. The unity thesis, to put it broadly, is the claim that the practical virtues (that is, virtues concerning action) are unified. The thesis says the virtues have some sort of harmonious connection with one another, rather than being mutually independent (lacking any connection whatsoever) or mutually incompatible (having a discordant connection). The thesis is of crucial significance for virtue ethics. For, if the virtues are mutually incompatible, they could not be reliably used to assess or guide actions. Similarly, an agent could not develop a virtuous character if the virtues are mutually

incompatible. In other words, if the practical virtues are disunified, they could lead one astray in action or in character development. Not surprisingly, then, virtue ethicists in the ancient Greco-Roman as well as medieval Christian traditions endorsed some form of the unity thesis.

There are two *prima facie* plausible versions of the unity thesis. It will be helpful to briefly consider these two versions of the thesis in order to develop an account of the unity of the virtues in which the notion of cardinal virtue plays a crucial role in unifying the virtues.

5.1.1 Partial and Full Unity of the Virtues

The first version of the unity thesis says that *all* the practical virtues are unified. Call this the full unity thesis. The most influential full unity thesis is the Aristotelian claim that all the virtues are unified by *phronesis*, or practical wisdom. Roughly, practical wisdom is the disposition to be appropriately sensitive to one's reasons for action, and thus act rightly, in any given circumstance. The Aristotelian claim is that to have any practical virtue at all one must have practical wisdom. The idea is that to have the virtue of generosity or honesty, say, one should know when and how to be generous or honest, and this requires practical wisdom. To have practical wisdom, in turn, is to have all the practical virtues because to have practical wisdom one must give appropriate consideration to all the relevant practical virtues in any given circumstance.

The second version of the unity thesis is the claim that the virtues are unified but only to limited degree. Call this the partial unity thesis. Neera Badhwar defends a version of the partial unity thesis (Badhwar 1996). On Badhwar's partial unity view, the virtues

are united within a particular domain, but not *across* domains. For instance, Mary could be generous towards her friends but lack generosity towards strangers. Hence, on her view there is *partial*, rather than *full*, unity of the virtues. Here, a domain is any aspect of an agent's life that is of practical significance that can be (i) psychologically isolated from other practical concerns, and (ii) is important enough to be the field of a virtue.[16] Domains are thus individuated partly by psychological considerations and partly by normative considerations.

Badhwar thus endorses a more modest version of the unity thesis compared to the Aristotelian version insofar as the scope of unity of the virtues is not global, encompassing practical normativity in its entirety. Like Badhwar, I endorse a version of the partial unity thesis. I leave open the possibility of the full unity of the virtues, but I am *not* committed to the fully unity thesis. An important difference, though, is that on my view virtues are united within any given *normative* domain. In other words, psychological considerations do not play any role in individuating the domains relevant for unity of the virtues. Rather, the relevant domains are individuated completely by normative considerations on my view.

The reason for the difference between Badhwar's partial unity thesis and mine lies in a fundamental difference in how we interpret the unity thesis. Badhwar interprets the unity thesis as an *empirical* claim that is true of any given (human) agent. However, I interpret the unity thesis as a *theoretical* claim that is true of any (satisfactory) theoretical model of the virtues. As such, this theoretical claim does not commit me to any empirical

[16] To be clear, Badhwar does not invoke a target centered conception of virtue. The notion of a virtue's field, however, is a useful way to present her view.

claims. For, it is possible that a satisfactory theoretical model of the virtues is sufficiently idealized such that what the model says is not true of any agent in the actual world. Such a model might still be useful insofar it does explanatory work and provides agents with an appropriate ideal to strive for.

5.1.2 Cardinal Virtue

Let us now consider my version of the unity thesis and the role that cardinal virtues play in unifying the virtues in their domain. My proposed version of the unity thesis is that any satisfactory theoretical model of the virtues of a normative domain must include at least one *cardinal* (domain-specific) virtue. Here, a cardinal virtue is a virtue that grounds another virtue in its normative domain and is not itself grounded by any virtue in its normative domain. Further, for a normative domain with more than one cardinal virtue, the cardinal virtues must be unified on any satisfactory theoretical model of the virtues of the normative domain. Let me explain.

It will be helpful to begin with a reminder of what I mean by a "normative domain". A (practical) *normative domain* takes as inputs certain facts about an act (or a disposition to act) and generates evaluative (e.g., good, bad) and deontic (e.g., required, permissible) outputs about the act (or the disposition to act). For instance, the fact that one might harm an innocent person if one were to drive drunk gives one a *moral* reason not to drive while drunk. Similarly, the fact that one might harm oneself or damage one's car gives one a *prudential* reason not to drive while drunk. While more could be said about individuating normative domains, for now I assume that morality and prudence are two distinct normative domains relevant for practical normativity.

Consider, now, the notion of a cardinal domain-relative virtue. Plausibly, the virtues of any given normative domain have a *hierarchical* structure, rather than a flat structure. That is, some virtues are especially important for any given normative domain. For instance, justice is plausibly an especially important moral virtue. Let us call such virtues *cardinal* virtues, following the traditional use of the term among virtue ethicists.[17] In what sense are cardinal virtues especially important? A plausible way to flesh out their importance is in terms of asymmetric relations of grounding — a cardinal virtue grounds some non-cardinal virtues in its domain and is not itself grounded by any virtue in its domain. That is, the (domain-relative) excellence of a non-cardinal virtue is derived from the (domain-relative) excellence of the cardinal virtue. For instance, the moral excellence of the disposition to keep one's promises is plausibly derived from the moral excellence of the disposition to be honest in one's dealings with others. Further, a cardinal virtue may be either the *complete ground* of a non-cardinal virtue or only a *partial ground*, depending on whether the cardinal virtue completely or only partially explains the excellence of the non-cardinal virtue.

When does a cardinal virtue ground a non-cardinal virtue in its domain? A necessary condition is that the field of the cardinal virtue must include the field of the non-cardinal virtue. Recall that the *field* of a virtue consists of the elements that are the sphere of concern of the virtue. If the field of a cardinal virtue CV does not include the field of a non-cardinal virtue NCV, then NCV could not possibly be grounded by CV since the sphere of concern of NCV is simply outside the sphere of concern of CV.

[17] I am not thereby endorsing the common list of cardinal virtues: phronesis, justice, temperance, and courage.

However, the scope of field is not by itself sufficient for CV to ground NCV. CV should also *explain* why NCV is a (domain-relative) virtue. That is, CV must provide the right kind of reasons for why NCV is a (domain-relative) virtue.

Importantly, a cardinal virtue helps *unify* the non-cardinal virtues it grounds. When the virtues of a normative domain are *disunified*, for some act (or disposition to act), the domain-relative evaluative and deontic properties of the act (or disposition to act), are indeterminate. For instance, when morality is disunified, the moral permissibility of lying to the Nazi at the door about the Jewish children hiding in one's basement may be indeterminate (because of the conflict between compassion and honesty). When the virtues of a domain are disunified they can pull in different, perhaps even opposite, directions in some cases and the normative output of the domain are indeterminate as a result. What is needed for unification of the virtues, then, is a mechanism to adjudicate or balance between these different or conflicting pulls.

A cardinal virtue helps unify the non-cardinal virtues that it grounds by providing them with an overarching aim (the target of the cardinal virtue) which allows us to rank the importance of the non-cardinal virtues and adjudicate in cases of conflict. Further, what is of particular interest for our purposes is the unity of *all* the virtues of a normative domain. Let us say that the virtues of a normative domain are *completely unified* just in case for any (disposition to) act, the domain-relative evaluative and deontic properties of the (disposition to) act are determinate. The simplest way such complete unity might come about is if there is a single cardinal virtue that grounds all the other non-cardinal virtues of its domain. All else being equal, a theoretical model of the virtues of a domain with *fewer* cardinal virtues is preferable because of the *simplicity* of the model. As such,

it would be methodologically best to begin with a single cardinal virtue for any given domain and only add additional cardinal virtues if necessary for explanatory or extensional adequacy. However, if any given domain has more than one cardinal virtue, the cardinal virtues of the domain must themselves be united in order for the normative domain to be completely unified.

In the next section, I sketch an account on which there is a single cardinal moral virtue that grounds all the other non-cardinal moral virtues. Before sketching the account, I should clarify that this account of unity of the moral virtues, unlike the Aristotelian view, requires no commitment to a *comprehensive* view of what it is to live a good life or to act well. It does, however, require commitment to the existence of distinct normative domains like morality and prudence. On this account of the unity of the virtues, settling questions about what it is to live a good life or to act well (all-things-considered) will involve some sensitivity to both moral and prudential considerations (as well as any other relevant considerations). But it leaves open the issue of these questions should be settled.

5.2 Universal Love as a Cardinal Moral Virtue

I now want to argue that the disposition to love universally is the sole cardinal moral virtue that grounds all the other (non-cardinal) moral virtues. My argument for this thesis is as follows. I start by showing that all the commonly recognized moral virtues *can* be completely grounded by universal love. Further, given the methodological assumption that, all else being equal, positing fewer cardinal virtues is preferable to positing more, I conclude that universal love is the sole cardinal moral virtue. The argument is an inference to the best explanation — if the resulting model of the moral virtues turns out to be explanatorily or extensionally inadequate, we would have to revise or reject the thesis.

In order to flesh out the argument, it is crucial to briefly recall the target-centered conception of virtues with which we are working. The *field* of a virtue consists of the elements that are the sphere of concern of the virtue. For instance, the field of benevolence is others' well-being, while the field of temperance is one's bodily pleasures. A *virtue* is a disposition to respond to elements within its field(s) in an excellent way, relative to a normative domain. For instance, benevolence is a morally excellent disposition to respond to elements of others' well-being. The *target* of a virtue is what the virtue is aimed at within its field. For instance, the target of benevolence is the promotion of others' well-being. A virtue's field thus specifies *what* the agent should respond to, while the virtue's target specifies *how* the agent should respond. Of particular relevance for our purposes is the moral virtue of universal love. The field of universal love is the well-being of all beings with moral standing who stand to be affected, in one way or another, by one's agency. The target of universal love is the promotion of the sum of the weighted, relevant interests of all beings with moral standing that stand to be affected by one's agency.

For (the moral virtue of) universal love to ground all the other moral virtues, its field must be sufficiently wide in scope to include the elements that are in the sphere of concern of all the other moral virtues. Further, universal love must also explain the virtuousness of all the other moral virtues. Let us now consider the common moral virtues to see whether this is the case. We shall consider six commonly recognized moral virtues: (i) benevolence, (ii) compassion, (iii) courage, (iv) honesty, (v) respect, and (vi) justice. I should clarify at the outset that individuation of the virtues is a matter of

ongoing research, and so the list is rather tentative, though I take the list to be plausible enough for our present purposes.

5.2.1 Benevolence

Benevolence is the disposition to promote others' well-being. The field of benevolence (others' well-being) is thus straightforwardly included in the field of universal love. When benevolence is limited in any number of ways — one might be benevolent only towards one's family or friends, or one's local or national community — the field of benevolence will be narrower than that of universal love. However, when benevolence is not limited in any of these respects but has all beings with moral standing as its object, the field of such *universal* benevolence is coextensive with that of universal love.

Further, universal love also explains why benevolence is a virtue. For, even in its limited forms benevolence is concerned with promoting the well-being of beings with moral standing and thus the agent closer to the ideal of universal love.

Indeed, I am happy to say that universal benevolence just *is* universal love as long as the target of universal benevolence is the aggregation of *relevant* interests, rather than a thoroughgoing aggregation of *all* interests.

5.2.2 Compassion

Compassion is the disposition to alleviate others' suffering. The field of compassion, then, consists of the negative elements of others' well-being. I have been using "well-being" in the broad sense that refers to how someone's life goes for them overall — the balance of the positive and negative elements of well-being in their life. The virtue of compassion is plausibly concerned with the negative elements of others' well-being. For

instance, compassion would dispose one to notice the starving person on the street and respond by helping them to feed themself.

Like benevolence, one's compassion may be limited in any number of ways. When one's compassion is not restricted, one would be concerned with the suffering of any being with moral standing. Even universal compassion, however, is concerned only with the *negative* elements of the well-being of all beings with moral standing, and not with the *positive* elements of their well-being. The field of compassion, then, is not co-extensive with that of universal love. The field of compassion is nonetheless included in the field of universal love since universal love is concerned with both the positive and negative elements of the well-being of all beings with moral standing.

Universal love also explains why compassion is a moral virtue, since an important part of being disposed to promote others' well-being is to be disposed to alleviate their suffering.

5.2.3 Courage

Courage is the disposition to protect others from dangers to their well-being, even when one's own well-being may be adversely affected.[18] The field of courage, then, are dangers to others' well-being. By a "danger" I mean anything that is likely to result in reduction of well-being over time, if not averted. Of course, some reductions in well-being at time *t* may be necessary in order to promote well-being at a later time *t'*. Such temporary reductions in well-being are not dangers, because they would not result in reduced well-being over time. For instance, quitting smoking might not be pleasant, but it plausibly

[18] Courage is plausibly also a prudential virtue insofar as it is a disposition to protect oneself from dangers to one's own well-being. My focus, however, is on courage as a moral virtue.

increases one's well-being over time. Quitting smoking, then, is not a danger though it results in a temporary reduction in one's well-being. Lung cancer, on the other hand, is a danger because it would result in reduction of well-being one's well-being over time.

How is the field of courage related to that of universal love? Here the relationship is not as straightforward as benevolence or compassion. For, a *danger* to someone's well-being is not *constitutive* of their well-being, rather it bears a *causal* relationship with someone's well-being. For instance, smoking is not by itself constitutive of reduction in one's well-being but rather causes a reduction in one's well-being. Nonetheless, the field of universal love should plausibly be understood as including not just the (positive and negative) constitutive elements of well-being, but also the *causes* of well-being. Here is some motivation for this claim. Consider love for someone as a virtuous disposition toward them. As we have discussed earlier, love for someone is the desire for the promotion of their well-being for their own sake. What, exactly, is the sphere of concern of such a desire? Surely it is concerned with the constitutive elements of their well-being, like pleasure. But plausibly such a desire must also be concerned with *how* to bring about their well-being. That is, the intrinsic desire to promote someone's well-being grounds an *instrumental* desire for things that promote their well-being, like drinking coffee or eating chocolate. The same considerations hold for universal love — the elements that are the sphere of concern of universal love are plausibly not just the constitutive elements of others' well-being but also the elements that are causally relevant for their well-being.

Further, if the field of universal love includes elements that are causally relevant for others' well-being, it will also include dangers. For, a danger is something that would cause a decrease in someone's well-being over time, if not averted. So, a danger bears a

causal relationship with a possible reduction in someone's well-being over time. Moreover, a danger can often be averted by one's agency. That is, one can do something to prevent the danger from being causally efficacious in reducing their well-being. Insofar as dangers bear this causal relationship with the well-being of beings with moral standing who stand to be affected by one's agency, the field of universal love includes the field of courage. Universal love thus also helps explain why courage is a moral virtue. The virtuousness of courage, from the moral point of view, consists in the disposition to avert dangers to others' well-being.

We have so far considered benevolence, compassion, and courage, and seen that the field of universal love plausibly includes the field of these moral virtues. These three moral virtues are concerned with *all* beings with moral standing as their object. However, the next three moral virtues we shall consider — honesty, respect, and justice — are concerned, at least primarily, only with those beings with moral standing that are also rational agents. While these latter moral virtues could be also applied to, say, children or (non-human) animals, I shall focus primarily on how they apply to rational agents.

5.2.4 Honesty

Honesty is the disposition to be truthful in interpersonal communication with other rational agents.[19] The field of honesty, then, is interpersonal communication with other rational agents.[20] Now, recall that on the account of well-being we are working with, knowledge is a constitutive element of the well-being of rational agents (chapter 2).

[19] Honesty, like courage, is also plausibly a prudential value.
[20] One might plausibly think that honesty is not just narrowly concerned with interpersonal communication but with interpersonal relationships more broadly.

Further, knowledge is not exclusively, or even primarily, an individualistic endeavor. Given our cognitive limitations as human beings, we often have to rely on others in order to gain knowledge. The social dimension of knowledge-gathering is most pertinent when it comes to gaining knowledge about other people's mental states. Interpersonal communication is also relevant for gaining knowledge through testimony about a wide variety of subjects — from current news to the natural and social sciences.

Interpersonal communication, then, is causally relevant for knowledge. Interpersonal communication can lead to gaining knowledge, or it can lead of deception. Knowledge, in turn, is a constitutive element of the well-being of rational agents. So interpersonal communication can lead to either increase or decrease in well-being (with respect to knowledge). Similar to courage, then, the field of honesty is included in the field of universal love insofar as interpersonal communication bears a causal relationship to the well-being of rational agents. The virtuousness of honesty, from the moral point of view, consists in the disposition to promote their well-being by being truthful in interpersonal communication.

5.2.5 Respect

Respect is the disposition to honor the autonomous choices of other rational agents.[21] The field of respect consists of the autonomous choices of other rational agents. Recall that on the account of well-being we are working with, autonomy is a constitutive element of the well-being of rational agents (chapter 2). The autonomous choices of other rational agents are thus straightforwardly included in the field of universal love. For, universal

[21] I am following the influential Kantian understanding of respect. There are, no doubt, other plausible ways of understanding respect.

love is concerned with the constitutive elements of well-being of beings with moral standing, which includes rational agents. So, the field of universal love includes, among other things, the autonomous choices of rational agents.

Further, universal love also explains why respect is a virtue. Part of promoting the well-being of rational agents is honoring their autonomous choices because their autonomy is an element of their well-being.[22]

5.2.6 Justice

Justice is the disposition to honor the moral claims of other rational agents.[23] The field of justice consists of the moral claims of other rational agents. As I understand it, justice is concerned negatively with ensuring that one does not wrong others by not honoring their claims, and positively with ensuring that one gives others their due by honoring their claims. To consider how moral claims are related to the field of universal love, we will first need to clarify the nature and grounds of moral claims.

What is a moral claim? For *A* to have a claim to *X* just is for *A* to have the authority to make demands against other agents concerning *X*. For instance, for Hannah to have a claim to bodily autonomy just is for Hannah to have the authority to make demands against other agents concerning the use of her body. This means, among other things, that she has the authority to demand that someone not touch her body without her consent. More generally, rational agents may have promissory and contractual claims,

[22] To be clear, this does not mean that one should *always* respect others' autonomous choices, as I discuss in the next chapter.

[23] As I clarified above, the virtue of justice might plausibly be thought to apply to children and (non-human) animals as well. However, I will focus on how it applies to rational agents, setting aside if and how it might be extended to children and (non-human) animals.

claims against harm, and claims to benefits. Further, if *A* has a moral claim to *X*, then *A* is *wronged* by other moral agents when they do not honor *A*'s claim to *X*. For instance, someone who touches Hannah's body without her consent wrongs her by not honoring her claim to bodily autonomy. This raises the further question of what is it to wrong someone? Among other things, being wronged gives one the authority to demand rectification from the wrongdoer, as well as the authority to blame and forgive the wrongdoer.

What are the grounds of moral claims? There are three possible views here. The first view would be that moral claims are basic, not grounded in anything more fundamental. On this view, a claim to bodily autonomy, for instance, is a basic claim not grounded in anything else. The second view would be that claims are grounded in the claimholder's rational will. The third view would be that moral claims are grounded in interests — one's claims depend somehow on one's interests. On this third view, a claim to bodily autonomy is grounded in one's interest in control over one's body. That is, one has a claim to bodily autonomy in virtue of having an interest in control over one's body. I prefer the third view and shall simply assume it as a working assumption. However, I do not make any assumptions about *how* claims are grounded in interests.

This is, of course, a very bare sketch of an account of moral claims. The important point for our present purposes is that if the moral claims of rational agents are grounded in their interests, then the field of universal love will include moral claims. This is because the field of universal love includes the constitutive elements of well-being of rational agents which in turn ground their moral claims. Universal love thus also explains why justice is a virtue. For, part of promoting the well-being of rational agents is

honoring their moral claims, which protect and promote the elements of their well-being like their bodily autonomy.

Let us take stock. We have just considered six commonly recognized moral virtues — benevolence, compassion, courage, honesty, respect, and justice — and concluded that the field of universal love is wide enough to include the elements that these virtues are concerned. Universal love also plausibly explains why these are moral virtues. Thus, the moral virtue of universal love (i.e., the disposition to love universally) may plausibly be understood as the ground of all the other moral virtues. In other words, all the other moral virtues may plausibly be understood as different aspects of universal love.

The discussion so far shows that universal love is a cardinal moral virtue that grounds all the commonly recognized moral virtues. Further, given the methodological assumption that, all else being equal, positing fewer cardinal virtues is preferable, we may tentatively conclude that universal love is the sole cardinal moral virtue. The conclusion is tentative because its plausibility depends on the explanatory and extensional adequacy of the resulting account of the moral virtues. To that end, I develop an account of moral reasons for action in the next section and shall defend its plausibility in the subsequent chapter.[24]

[24] To be clear, an exhaustive account of morality will not just provide an account of moral reasons for action but also have something to say about other aspects of morality like moral character and social institutions.

5.3 Universal Love and Moral Reasons for Action

The disposition to love universally, by virtue of being the sole cardinal moral virtue, also unifies the moral virtues by providing an overarching aim that helps adjudicate conflicts between the other (non-cardinal) moral virtues. As a result, universal love is conclusively relevant for determining moral reasons for action.

Let me start by clarifying that I am assuming a virtue ethical framework about reasons for action. Recall the virtue-based account of reasons for action sketched in the previous chapter. On this framework, (domain-relative) virtues are the grounds of (domain-relative) reasons for actions. On the target-centered conception of virtues, an act *A* is virtuous with respect to some virtue *V* to the degree that *A* achieves the target of *V*. An act's being virtuous (to whatever degree) in turn provides *pro tanto* (domain-relative) reason for performing the act (to the degree of its virtuousness). For instance, an act is virtuous with respect to benevolence to the degree that it successfully promotes others' well-being in a given circumstance, which in turn provides *pro tanto* moral reason to perform the act. On this framework, moral reasons for action are completely determined by the moral virtues.

Now, as discussed in the previous section, disunity arises between the virtues of a given normative domain when they pull in different directions in particular circumstances, resulting in indeterminacy of domain-relative reasons for action. A cardinal virtue helps unify the (non-cardinal) virtues it grounds by providing an overarching aim (the target of the cardinal virtue) which allows us to rank the importance of the non-cardinal virtues and adjudicate in cases of conflict. When a normative domain has just one cardinal virtue, the cardinal virtue provides conclusive domain-relative

reason for action. Thus, in the case of morality, universal love is conclusively relevant for determining moral reasons for action. Universal love provides an overarching aim (the target of universal love) which allows us to rank the importance of particular moral virtues and adjudicate in cases of conflict. For instance, consider the conflict between honesty and compassion when confronted with the Nazi at one's door asking about the Jewish children hiding in one's basement. Universal love allows us to adjudicate the conflict by appealing to the overarching aim of promotion of the sum of the weighted, relevant interests of all beings with moral standing that stand to be affected by one's agency. Arguably, in this case the Nazi's interest in not being deceived is not even *relevant* to the children's interest in not being killed or sent to a concentration camp. At least in this case, then, the conflict is easily resolved by appealing to the target of universal love.

An act's being virtuous with respect to universal love thus provides (conclusive) moral reason for performing the act (to the degree of its virtuousness). Virtuousness with respect to universal love thus grounds moral reasons for actions. This allows us to define the moral properties of acts by appealing to the target of universal love. Recall that the target of universal love is the promotion of the sum of the weighted, relevant interests of all beings with moral standing that stand to be affected by one's agency. Setting the sum of the strength-weighted, relevant interests if one were to refrain from action as the baseline, we can define the evaluative moral properties of actions as follows:

(i) Morally bad act = an act that *decreases* the sum of the strength-weighted satisfaction of relevant interests (relative to inaction).

(ii) Morally neutral act = an act that *does not change* the sum of the strength-weighted satisfaction of relevant interests (relative to inaction).

(iii) Morally good act = an act that *increases* the sum of the strength-weighted satisfaction of relevant interests (relative to inaction).

Further, the deontic moral properties of actions can be defined in terms of their evaluative moral properties. I merely discuss three possible views without committing to any of them:

(iv) Morally permissible act (three possible views):

a. Permissive view: Morally permissible act = an act that is not bad (i.e., it is neutral or good).

b. Moderate view: Morally permissible act = an act that is at least good enough (on the relevant measure of "good enough").

c. Demanding view: Morally permissible act = an act that is maximally good (relative to all other feasible good acts).

6 In Defense of Morality as Universal Love

In the previous chapter, I sketched a virtue ethical moral theory on which the disposition to love universally is the sole cardinal moral virtue. That is, on the proposed theory, all the (other) moral virtues are grounded in (the moral virtue of) universal love, but universal love is not itself grounded in any other moral virtue. Universal love is thus the sole foundational moral virtue on this theory. In this chapter, I want to defend this theory by addressing some important objections.

Here is a brief overview of the objections I will consider below. The first objection is that universal love is not a cardinal moral virtue because one can have *excessive* love for all beings with moral standing. The objection is that any character trait that can be had in excess is not a virtue, let alone a *cardinal* virtue. The second objection is that respect (primarily for rational agents, but also perhaps for other beings) is a moral virtue but it cannot be grounded in universal love because love and respect are fundamentally different kinds of attitudes. The third objection is a similar worry that justice cannot be grounded in universal love because love and justice are, at least on occasion, opposed to each other. Lastly, the fourth objection is that my account of universal love qua a moral virtue is mistaken insofar as it allows for some interpersonal aggregation. The objection is that universal love is incompatible with *any* interpersonal aggregation because love treats the well-being of the beloved as non-fungible.

6.1 Is Universal Love a (Cardinal) Virtue?

The first objection I want to consider is that universal love is not a cardinal moral virtue. I have been arguing that the disposition to love universally is the foundational moral virtue, but one might question the status of universal love as a cardinal moral virtue.

Why might one question the status of universal love as a cardinal moral virtue? To begin with, universal love does not appear on a traditional list of the cardinal virtues popular in the ancient Greco-Roman and medieval Christian traditions.[25] A bare appeal to tradition, however, is not in itself strong reason to question the plausibility of a non-traditional view. Moreover, as I briefly mentioned in the introduction to the dissertation, there are ethical traditions — for instance, Buddhist and Christian — that give primacy to universal love. So, at best, whether or not universal love is a traditional cardinal virtue depends on the ethical tradition(s) under consideration.

There is, however, a recent defense of the traditional list of the cardinal virtues that warrants consideration. Paul Bloomfield (2022) has recently argued that some surprising conclusions follow from the conceptual claim that virtues are *excellences* of character. Let us briefly consider these alleged conclusions and their relevance for the cardinality of universal love.

First, since the virtues are excellences of character, they can never lead an agent astray in her practical deliberation and agency. That is, virtues guarantee right action because they are excellences of character. For, if a character trait sometimes leads an agent to do the wrong action it would not be an *excellence* of character. This is because the *proper function* of a virtue is to reliably guide an agent on how to act in the virtue's characteristic sphere of concern. For instance, the virtue of courage reliably guides an agent on how to deal with dangerous situations.

[25] Prudence, fortitude, temperance, and justice being the four cardinal virtues according to this traditional list.

Second, since virtues are excellences, an agent cannot be excessively virtuous, because being *too excellent* is an incoherent notion. This further entails that any character trait that *can* be had in excess is not a virtue, because anything that can be over-developed cannot be an excellence. The thought is that anything that has a proper function can either fulfill its function or fail to do so, but it cannot function *too* well. As a result, just as one cannot be too healthy, one cannot be too virtuous either.

Bloomfield's radical conclusion, then, is that the application of virtue terms must be restricted to character traits that cannot be over-developed. Indeed, he suggests that failure to restrict the application of virtue terms threatens to undermine virtue ethics since without such a restriction the virtues can lead agents astray in their practical deliberation and agency. This undermines the virtue ethical emphasis on virtues as reliable guides for how one should act. Further, Bloomfield argues that such a restriction on the application of virtue terms leads to the conclusion that only the traditional cardinal virtues — justice, prudence, courage, and temperance — are virtues. This is because, he argues, one cannot have these traits in excess. However, on his view, character traits like honesty or generosity are not virtues, because one can be too honest or too generous.

Bloomfield's argument is an important objection to the view under consideration because if successful it would rule out the possibility of universal love being a virtue, let alone a *cardinal* moral virtue. Before fleshing out my response to his argument, let us start by clarifying what it would be to have the character trait of universal love in "excess". Recall that the disposition to love universally is the disposition to desire the well-being of all beings with moral standing (that one veridically believes to exist) for their own sake. An *excessive* desire for the well-being of all beings with moral standing

would be to desire the well-being of all (other) beings with moral standing *more* than one's desire for one's own well-being. Such an agent might desire the well-being of her family and friends more than her own well-being. She might even desire the well-being of strangers more than she desires her own well-being. In general, such an agent will be disposed to make sacrifices in order to benefit others. Call such sorts of desires *saintly* dispositions.

Saintly dispositions are plausibly "excessive" forms of universal love. Does it follow that they are not virtuous character traits? I suggest not. Let me briefly summarize my response before fleshing it out in detail. My response to Bloomfield's argument is two-fold. First, any particular virtue only provides *pro tanto* reasons for action. It follows that that any individual virtue can, by itself, lead an agent astray. However, this is consistent with all the virtues (collectively) providing an agent with reliable normative guidance. Second, we must make a distinction between moral and non-moral reasons for action. The judgement that universal love can be "excessive" is plausibly made from a non-moral, prudential perspective, rather than from the moral point of view. In other words, saintly forms of universal love may simply manifest the tension between moral and non-moral excellence. I do not take a stance on how, or even if, such conflicts should be resolved. Collectively, these two responses uphold the status of universal love as cardinal moral virtue against Bloomfield's argument. Let us now consider each response in more detail.

First, any particular virtue only provides *pro tanto* reasons for action. Bloomfield argues that the proper function of a virtue is to reliably guide an agent on how to act in the virtue's characteristic sphere of concern. I agree. But it does not follow that a virtue,

by itself, can never lead an agent astray in her practical deliberation and agency. That only follows if we take each virtue to provide *conclusive* reasons for action. However, a more plausible view is that any given virtue only provides *pro tanto* reasons for action. On a virtue ethical framework, the virtues collectively provide reliable normative guidance for how to act, but one need not expect any particular virtue, by itself, to provide reliable normative guidance. I agree with Bloomfield that each virtue has its characteristic sphere of concern in which it provides reliable normative guidance. But in many, perhaps most, situations there will be many virtues relevant to the circumstance because their spheres of concern overlap under the given circumstances. For instance, an action's being honest or generous only provides *pro tanto* reason for performing the action. An agent may have strong *pro tanto* reason to be brutally honest (because it would be honest), but stronger *pro tanto* reason *not* to be brutally honest (because it would be brutal). Of course, it will require higher order (cardinal) virtues, as well as practical wisdom, in order to discern the overall balance of reasons.

Second, we must make a distinction between moral and non-moral reasons for action. I must start by acknowledging that this is not a distinction that virtue ethicists commonly make. Nonetheless, it is a distinction that many other moral philosophers make, and it is helpful for virtue ethics as well. How exactly the distinction should be made is relatively controversial, but for our purposes it should suffice to make a distinction between moral reasons for action that concern others' well-being and prudential reasons for action that concern the agent's own well-being.[26] What is crucial for our purposes is that virtues are defined by reference to a particular domain. That is, a

[26] This is not meant to be exhaustive, and I leave open the possibility of other reasons for actions.

moral virtue is a character trait that is excellent from the moral point of view. Similarly, a prudential virtue is a character trait that is excellent from the prudential point of view.[27] The reason why this distinction matters is that the judgement that universal love can be excessive is made from a non-moral, prudential perspective — not from the moral point of view. To be clear, some character traits can be excessive from within a given normative perspective, say morality. For instance, one can be excessively honest insofar as one is disposed to tell the truth even doing so is not morally desirable (all-things-considered). My present suggestion, however, is that saintly forms of universal love are judged as "excessive" from a prudential point of view because of the sacrifice of the agent's own interests. But it is plausible that saintly forms of universal love are morally admirable. Indeed, saints are usually thought of moral exemplars precisely because they embody such morally admirable character traits. These intense forms of universal love only highlight the possibility of moral excellence being in tension with non-moral excellence.

Collectively, these two responses uphold the status of universal love as cardinal moral virtue against Bloomfield's arguments. The upshot is that the disposition to love universally, even in its intense saintly forms, is morally excellent and therefore a moral virtue. Importantly, my view does not require that an agent cultivate saintly forms of universal love. Saintly forms of universal love are certainly morally admirable, but they are by no means morally required. On my view, agents ought to cultivate a desire for the well-being of all beings with moral standing for their own sake. However, this is

[27] Of course there can be an overlap. Some character traits might be excellent from both the moral and the prudential points of view.

compatible with a stronger desire for the well-being of some (relative to others), and with a stronger desire for one's own well-being (relative to all others).

6.2 Love and Respect

The second objection I want to consider is that respect cannot be grounded in universal love because love and respect are fundamentally distinct kinds of attitudes. Many philosophers who recognize the moral significance of love nonetheless regard love to be distinct from, and indeed in some ways *opposed* to, respect. The thought is that even if love is relevant for morality, love is not *all* that is relevant for morality. Further, the alleged tension between love and respect is morally desirable on this view. For, the thought goes, left to its own devices love would lead one to act in ways that are intrusive to the beloved. Love, then, must be balanced by respect on this view. As a representative of this view, consider this passage from Kant's *Doctrine of Virtue*:

> "In speaking of laws of duty (not laws of nature) and, among these, of laws for men's external relations with one another, we consider ourselves in a moral (intelligible) world where, by analogy with the physical world, attraction and repulsion bind together rational beings (on earth). The principle of mutual love admonishes men constantly to come closer to one another; that of the respect they owe one another, to keep themselves at a distance from one another; and should one of these great moral forces fail, "then nothingness (immorality), with gaping throat, would drink up the whole kingdom of (moral) beings like a drop of water".
>
> (*The Metaphysics of Morals*, 449)

In this passage Kant paints a stark picture of love and respect as mutually opposing moral forces. To use Kant's analogy, love is the moral force of attraction that compels us to come closer to one another while respect is the moral force of repulsion that compels us to keep our distance from one another. Plausibly, friendships and romantic relationships are forms of coming close, while strictly professional relationships like the physician-

patient or student-teacher relationships are forms of keeping one's distance. Kant's thought is that the moral life needs both of these ways of relating to others.

Kant's picture of the mutual opposition between love and respect has been rather influential. Christine Swanton, for instance, defends Kant's picture of *love* as *coming close* and *respect* as *keeping distance* (Swanton 2003: 104-110). She argues that Kant's analysis helps illuminate the proper limits of love. She argues that love that is not limited by respect can lead to a host of problematic behaviors — benevolence can turn into disrespectful paternalism, loving parental concern can turn into disrespectful interference, romantic love can turn into possessiveness, etc. (Swanton 2003: 106). On her view, Kant's analysis of love and respect as distinct and opposing moral forces helps illuminate the role of love in the moral life.

This Kantian view of the relationship between love and respect is an important objection to the view I sketched in the previous chapter. On my view, respect is an *aspect* of love for persons[28]. Importantly, if the Kantian analysis is right — love and respect are *independent* and *opposing* disposition — then respect would not be grounded in love, rather it would be an independent moral virtue.

How do we adjudicate between these two views on the relation between love and respect? To adjudicate between these views, we need to first clarify the meanings of "love" and "respect". Different ways of understanding love and respect will no doubt lead

[28] To keep the discussion focused I will focus on respect for persons, i.e., rational agents, setting aside the issue of respect for children and animals.

to different analyses of how they are related. I want to argue that given my account of love and a plausible understanding of respect, respect just is an aspect of love.

Recall that on my view, love for someone is the desire for the promotion of their well-being for their own sake. In the last chapter, I defined respect as the disposition to honor the autonomous choices of other rational agents. So understood, respect is an aspect of love. For, part of what it is to love a *person*, i.e., a rational agent, is to desire the promotion of their autonomy. This is because autonomy is an element of a person's well-being. Let me elaborate.

We may understand autonomy as the capacity to engage in practical deliberation and realize one's intentions. For ease I shall refer to these two capacities as "rational capacities" in what follows. The capacity to engage in practical deliberation is the capacity to consider different options and one's reasons in favor of and against realizing those options. The output of practical deliberation is an intention to realize some action or state of affairs. The capacity to realize one's intentions is the capacity to reliably bring it about that the action or state of affairs obtains. Take a mundane example — I am deciding what to eat for lunch. I consider a few different options that I could cook, get takeout, or delivered. I decide that I want to get Mexican takeout. This is all part of practical deliberation. I then go to my favored Mexican restaurant and get a takeout. This is part of intention realization. Throughout this whole episode I am exercising my autonomy in a rather mundane scenario.

Promoting someone's autonomy could involve: (i) promoting the *development* of their rational capacities, (ii) honoring the *exercise* of their rational capacities, and (iii) promoting the *opportunity* to exercise their rational capacities. (i) would be primarily

relevant for children who are developing their rational capacities, though it would also be relevant for adults who have not developed their rational capacities because of social or psychological developmental impediments. (ii) would be allowing, or indeed encouraging, people to exercise their rational capacities, say by allowing or encouraging them to make important choices about their family, religion, or career. (iii) would be increasing the number of feasible options among which one could choose, say in choice of one's spouse, religion, or career. For our purposes, (ii) is the most relevant, since (i) is relevant primarily in child rearing while (iii) is determined primarily by the kind of society one lives in.

Respect for someone, in the context of interpersonal morality, just is the desire for the promotion of exercise of their rational capacities. In particular, this desire manifests as the *pro tanto* disposition to *honor* their autonomous choices. Consider how this works in a clinical context. A physician who wants to promote her patient's autonomy would provide her patient with adequate information about the different treatment options and allow her to make an informed choice. Further, she would then honor the patient's choice and use the treatment option the patient has chosen. Notice that the physician could be said to keep her distance from the patient insofar as she allows the patient to choose, instead of paternalistically picking a treatment option on the patient's behalf. Nonetheless, the physician is manifesting love for her patient insofar as she is concerned about promoting her patient's autonomy.

More generally, the desire to promote someone's autonomy manifests itself as a *pro tanto* disposition to honor their autonomous choices[29]. Respect, then, just is an aspect of love for other rational agents. Understanding respect as an aspect of love for persons allows us recast Swanton's worries about the limits of love as *internal* limitations within the moral virtue of love, rather than external limitations imposed by an independent and opposed moral virtue of respect. Consider again some of Swanton's worries about the limits of love. Plausibly, benevolence turns into disrespectful paternalism when the desire to promote someone's well-being at the expense of promoting their autonomy in fact reduces their overall well-being. That is, in such cases the reduction in well-being because of decrease in autonomy outweighs any potential increase in other elements of one's well-being. Similarly, loving parental concern turns into disrespectful interference when the parent's interference in the child's life in fact decreases the child's overall well-being by stunting the development of their rational capacities.[30] The other cases Swanton mentions are easily seen to be unloving. It is hard to see how possessiveness is loving since possessiveness is a result of concern with one's self-interested desires, rather than a desire for the good of the other.

A last clarificatory point. Even if respect for someone is the desire for the promotion of their autonomy, the question remains open: Under what circumstances is it justified to override someone's autonomous choices? In other words, under what circumstances is paternalism justified? My short answer is that paternalism is justified

[29] To be clear, it also manifests as the disposition to promote the development of their rational capacities and the disposition to promote the opportunities to exercise their rational capacities.
[30] The case of children is admittedly more complicated since they are only developing their rational capacities.

when it does, in fact, promote someone's overall well-being. Of course, this is rather vague and spelling out the answer in more detail will require a detailed account of well-being that could inform a relative weighing of the different elements of the well-being of rational agents. But that is a task for another time.

I conclude, then, contra Kant and Swanton, that respect is an aspect of love rather than an independent and opposed moral virtue.

6.3 Love and Justice

The third objection I want to consider is similar to the second. The objection is that justice cannot be grounded in universal love because love and justice can, on occasion, come into conflict. Let me start by motivating and fleshing out the objection before responding to it.

Given the nature of universal love and justice, it is *prima facie* plausible that they could come into conflict. In particular, there are at least two kinds of *prima facie* conflicts that may arise between justice and universal love. First, justice might prohibit the promotion of well-being in some cases. That is, in some cases promoting well-being would be unjust. For instance, if an agent could save the lives of two innocent people only by killing a third innocent person, presumably universal love would recommend, or perhaps even require, performing the unjust act. Second, justice might require, or at least recommend, the reduction of well-being in some cases. For instance, an agent who has committed severe crimes might deserve to be punished, which by definition would involve a reduction of their well-being.

Now, if these are genuine conflicts, rather than merely apparent conflicts, this would be a problem for morality as universal love. For, if justice is grounded in universal love, then justice would be in harmony with universal love. To see the point, consider an analogous grounding claim. Suppose that Sal's being in the USA is grounded in the fact that he is in Chicago. It is not possible that Sal could be in Chicago and yet fail to be in the USA (at the same time). On the other hand, Sal's wearing a red shirt is not grounded in the fact that he is in Chicago. It is perfectly possible for Sal to be in Chicago but fail to wear a red shirt because his being in Chicago and wearing a red shirt are independent states of affairs. Similarly, if there can be genuine conflicts between universal love and justice, it follows that they are two independent virtues, rather than justice being grounded in universal love as I have argued. To defend morality as universal love, then, we need to show that there is no genuine conflict between universal love and justice in the kinds of cases under consideration. Let us consider each kind of case in turn.

The first kind of case is where justice prohibits the promotion of well-being because promoting well-being would be unjust. These are cases where promotion of well-being comes at the cost of violating moral claims. To consider the example form earlier — perhaps an agent can save the lives of two innocent people only by killing a third innocent person. In this case, killing the innocent person would be unjust because it would violate their moral claim against being lethally harmed.

Now, a thoroughgoing aggregationist view would recommend, or perhaps even require, performing the unjust act in order to promote overall well-being. Recall that I have endorsed a limited aggregationist view on which the target of universal love is the promotion of the sum of the weighted, relevant interests of all beings with moral standing

that stand to be affected by one's agency. Recall further that one of the factors for determining the strength of interests is whether the agent is *causing* harm or merely *allowing* harm. Plausibly, when aggregating others' interests, all else being equal, a moral agent should give significantly greater weight to others' interest in not being harmed by her compared to others' interest in being rescued by her. While I am skeptical about the prospects of specifying the weights with precision, it is clear enough that someone's interest in not being killed by an agent is at least twice as strong as their interest in being rescued from death by the agent. Indeed, it seems plausible that relative weight is much greater than that. Importantly, if that is right, then universal would not recommend killing one innocent person to save the lives of two other innocent people.

To be clear, such a response leaves it open that there are some cases where the overall weighted well-being would be promoted by harming some individuals. However, I do not see this as a problematic implication. Indeed, in such cases I would say that it would not be *unjust* to harm the individuals. This is because in such cases the moral claims of the individuals harmed would be genuinely outweighed by the competing moral claims of other individuals. I should also clarify that on my view only *relevant* claims are aggregated. This means that some claims might not be relevant for aggregation. For instance, universal love does not recommend torturing one person for an hour to relieve the mild headaches of a thousand people for one hour. In other words, the justificatory threshold for inflicting harms will depend on a variety of factors like the kinds of harms others are being saved from, how many are being saved, etc.

The upshot is that universal love grounds *pro tanto* constraints, but not *absolute* constraints, against being harmed. Beings with moral standing have *pro tanto* moral

claims against being harmed, but these claims may sometimes be permissibly violated, without being extinguished altogether. In particular, when the overall weighted well-being would be promoted, moral claims may permissibly be violated.

Consider next, the second kind of case where justice might require, or at least recommend, the reduction of someone's well-being. One might think, for instance, that an agent who has committed severe crimes deserves to be punished. Such punishment would, by definition, involve a reduction of the agent's well-being and be at odds with love for the agent.

Let me quickly flag two responses to this issue that are misguided in my opinion, before sketching my preferred response. One misguided response to this issue would be to say that being punished is *good for* the agent punished and therefore an expression of love for them. Perhaps one might say that insofar as being punished is part of the agent's moral education, it is good for them. However, it is implausible that punishment could be good for an agent, since punishment by definition seeks to reduce their well-being. Another response would be reject the claim that anyone ever *deserves* to be punished. However, even a purely consequentialist justification for punishment would leave the problem of treating the agent in an unloving manner for the sake of the greater good. While, on the other hand, rejecting the institution of punishment altogether would plausibly lead to worse outcomes overall and thus be incompatible with universal love.

Here is a sketch of my preferred response. First of all, I do think agents who have committed severe crimes, like murder or rape, deserve to be punished. Further, being punished is not good for the one being punished, insofar as punishment reduces the agent's well-being by, at the very least, reducing their autonomy. Nonetheless, the

(weighted) harm done to the agent punished could plausibly be outweighed by the (weighted) benefit to the victim of the wrongdoing, as well as other members of society. There are three important claims to consider here. First, while committing severe crimes does not warrant cessation of love for the agent altogether, it may warrant a *weaker* love. That is, the fact that Bob is a serial murderer might be reason to give Bob's well-being lower weight. Second, the victim of a severe crime might have a moral claim that the perpetrator be punished. Or if the victim is deceased, the victim's family or their community at large might have a moral claim that the perpetrator be punished. Third, punishment might have overall beneficial consequences insofar as it reduces crime and therefore be conducive to promoting others' well-being. My claim, then, is that the (weighted) benefit to the victim of the wrongdoing and other members of society can plausibly outweigh the (weighted) harm done to the agent punished. If that's right, universal love may recommend, or even require, punishing agents who commit severe crimes.

To sum up: We considered two kinds of cases where there is an apparent conflict between justice and universal love — cases where justice might prohibit the promotion of well-being or justice might require, or at least recommend, the reduction of well-being. I have argued that these are not cases of genuine conflict between justice and universal love. On my view, these cases merely raise the issue of under what circumstances universal love recommends, allows, or requires harming some in order to benefit others. The demands of justice turn out to be aspects of universal love that, no doubt, need to be balanced with other aspects of universal love.

6.4 Universal Love and Interpersonal Aggregation

The last objection I want to consider is that my account of universal love as a moral virtue is mistaken insofar as it allows for interpersonal aggregation. The objection is that universal love is fundamentally anti-aggregative because love treats the beloved as non-fungible. Universal love, the objection goes, is incompatible with both (i) aggregating benefits across individuals, and (ii) aggregating harms to some with benefits to others. The objection, then, is that I have misunderstood the nature of universal love as a moral virtue, even if I am right about its centrality for moral theory within a virtue ethical framework.

It will be helpful to begin by reiterating how I understand the virtue of universal love. The target of universal love, on my view, is the promotion of the sum of the weighted, relevant interests of all beings with moral standing that stand to be affected by one's agency. This is a partial-aggregative view on which universal love justifies the interpersonal aggregation of *relevant* interests, but not that of irrelevant interests. Here relevance is a context-sensitive notion. An interest is relevant in any given circumstance if it has significance for moral deliberation in that circumstance. An interest's deliberative significance, in turn, depends on other interests at stake in the given circumstance. For instance, a minor headache for an hour is not relevant when someone's life is at stake, though losing all of one's limbs is relevant. Further, all the relevant interests, in any given circumstance, can be assigned weights and the target of universal love is to achieve a maximal sum of the weighted, relevant interests.[31]

[31] Of course, there might be many maximums that can be achieved. In the simplest case, one might achieve a maximum sum in many ways depending on the specific individuals one helps or harms.

The objection under consideration rejects this account of the virtue of universal love. The objection goes something like this: "Love for someone is the desire to promote their well-being for their own sake. Universal love, then, is the desire to promote the well-being of each being with moral standing for their own sake. This rules out decreasing the well-being of *any* being(s) with moral standing for the sake of increasing the well-being of others since doing so would *not* manifest love for the being(s) harmed. Nor is it consistent with universal love to add up the benefits to different beings with moral standing because it fails to manifest love for *each* individual. The target of universal love, instead, is to promote the well-being of any being with moral standing, without harming any being with moral standing."

One finds variants of this objection raised in the literature. For instance, Kieran Setiya (2014) argues that another's humanity justifies loving them, which in turn justifies acting with partiality towards them. A crucial implication of this view is that the numbers do not count — faced with the choice to save one or five, an agent is permitted to save the one because she is justified in loving the one and acting with partiality towards them. To be clear, Setiya does not explicitly discuss universal love, though his views suggest a fundamentally anti-aggregative picture of universal love.

More explicitly, Quinn White (2025) argues that universal love justifies an anti-aggregative ethic of respect.[32] On his view, universal love is a *moral ideal* which is too demanding for moral agents to manifest. Accordingly, manifesting universal love is supererogatory and what is *required*, instead, is a *minimal approximation* of universal

[32] I should add that, like Setiya, White's focus is on love for human beings. He does not include animals as objects of universal love in his discussion.

love. He goes on to suggest that the minimal required approximation of universal love just is universal respect. Part of universal respect, in turn, is to treat others as non-fungible, which rules out aggregating benefits across individuals or harming some in order to benefit others.

These are two concrete examples of the general objection that since love treats the beloved as non-fungible, universal love is *not* compatible with any kind of interpersonal aggregation. In response, I want to argue that while the anti-aggregative understanding of universal love takes seriously the separateness of the objects of universal love (i.e., beings with moral standing), it fails to adequately take into account the *universality* of universal love. The partial-aggregative understanding of universal love, on the other hand, adequately takes into account both (i) the separateness of the objects of universal love, and (ii) the universal scope of universal love. Let me explain.

The field of universal love is the well-being of all beings with moral standing who stand to be affected by one's agency. How should universal love guide an agent's moral deliberation? One aspect would be that the well-being of each individual would have intrinsic deliberative significance. However, another aspect would be that the well-being of *all* individuals has intrinsic deliberative significance. Both these aspects together justify, in some but not all cases, (i) aggregation of benefits across individuals, and (ii) aggregation of benefits to some with harms to others.[33]

[33] I add the qualifier "some" because aggregation is limited to relevant interests on my view. This distinguishes my view from thoroughgoing aggregationist views which do not restrict aggregation in any way.

6.4.1 Aggregation of Benefits

Consider first the aggregation of benefits across individuals. On my view, only *relevant* benefits are to be aggregated. As long as we restrict our attention to relevant benefits, the numbers should count on my view. Consider the case where one can only save lives of five innocent strangers or one innocent stranger, but not all six lives. On my view, universal requires one to save the five because doing so promotes the sum of the weighted, relevant interests. All six individuals have an interest in continued existence and given the circumstances their interests have equal weight (since they are all strangers and have equal moral status). Saving the one is impermissible in this case because it does not promote the sum of the weighted, relevant interests.

On the other hand, the anti-aggregationist view says that since each individual is non-fungible it is not compatible with universal *love* to aggregate each individual's interest in continued existence. In response, I want to say first of all that the anti-aggregationist view does not adequately take everyone's well-being into account. For, the view says that it is permissible to save the one. However, it is hard to see how one is taking the well-being of all six individuals seriously if one is willing to let five of them die in order to save one. The anti-aggregationist might say, in response, that her view takes everyone's well-being into consideration insofar it with gives each individual a fair chance to be saved. However, the willingness to let five die in order to save one only makes sense as a willingness to be arbitrarily partial to the one (as Setiya explicitly argues). I agree with the anti-aggregationist view that universal love requires a *pro tanto* desire to save each individual, but I disagree with the anti-aggregationist that, all things considered, universal love allows one to save the one at the expense of the five. The anti-

aggregationist view, at best, takes everyone's well-being into consideration to a *limited* extent.

At this point, the anti-aggregationist is apt to raise the worry that I have confused universal *love* with the *aggregative benevolence* of the sort that many consequentialists appeal to. But I want to argue that this is not the case. While my view does allow, and indeed requires, aggregation in some cases, the justification for doing so is not because it produces better consequences from some impartial standpoint. Rather, my view allows, or requires, aggregation because one ought to desire the well-being of each and every being with moral standing for their own sake. For instance, one ought to save the five rather than the one, because (i) one ought to desire the survival of all six individuals for their own sake, and (ii) one ought to act to realize as many of one's desires as possible. To be clear, in this case the desires ought to be of equal strength since there's no relevant difference between the six individuals. More generally, one ought to act to realize the maximal sum of strength-weighted desires for the well-being of beings with moral standing that stand to be affected by one's agency.

Further, aggregation is not *always* permissible on my view. Consequentialist views typically endorse thoroughgoing aggregation because the fundamental goal of moral agency, and perhaps practical rationality more generally, is to produce the most good. Such thoroughgoing aggregation entails that given a choice between saving a life and saving millions from a minor hour-long headache, it might be permissible to save millions from a minor headache. However, on my view such aggregation is impermissible because an interest in avoiding a minor hour-long headache is simply not relevant to an interest in continued existence. This is because the fundamental goal of

moral agency, on my view, is manifestation of love for all beings with moral standing, not the production of the most good.

6.4.2 Aggregation of Benefits with Harms

Consider next the aggregation of benefits to some with harms to others. Once again, giving everyone's well-being adequate deliberative significance does not rule out harming some in order to benefit others. Indeed, it *requires* the willingness to harm some in order to benefit others. To see this, consider the following. Suppose that, given the circumstances, Joe and Bob stand to be affected by Sally's agency. Suppose that Joe faces an imminent threat of death and, for whatever reason, the only way Sally can save his life is by harming Bob in some trivial manner — perhaps by pushing or pinching him. How would universal love guide Sally's deliberation? Surely, she ought to take both Joe's and Bob's well-being seriously for their own sakes. But suppose that Sally deliberates that she cannot inflict even trivial harm on Bob in order to save Joe's life. In so deliberating she is certainly taking Bob's well-being seriously, but it is hard to see how she is taking Joe's well-being seriously given that Joe has everything to lose in this case. It would seem that if Sally truly loves both Joe and Bob, she will open to inflicting at least trivial harms on Bob in order to save Joe's life.

More generally, universal love must be both universal *and* recognize the separateness and non-fungibility of beings with moral standing. The resulting view, as I have suggested earlier, is a partial-aggregative one that justifies the aggregation of *relevant* interests, but not that of irrelevant interests. For instance, in the case just considered above, Bob's interest in avoiding a trivial harm is simply not relevant to Joe's interest in continued existence. So universal love, I suggest, would justify inflicting

trivial harm on Bob in order to save Joe's life. However, if Joe's life could be saved only by *seriously* harming Bob, say by maiming him, plausibly universal love would *not* justify harming Bob. That is because Bob's interest in not being maimed, say, would be *relevant* to Joe's interest in continued existence. Further, Bob's interest would have greater weight than Joe's because Bob's interest is in not being harmed by Sally, while Joe's interest is in not being *allowed* by Sally to be harmed.

However, if the lives of high enough number of innocent persons, say a million, could be saved only by seriously harming one innocent person, then universal love would justify harming the one. One might object that in such cases the person harmed would not be recognized and loved as a distinct, non-fungible individual. But the alternative in such cases is to not recognize and love the other million individuals whose well-being is at stake. So, the choice is to either act unlovingly towards the one or to act unlovingly towards the other million individuals. I suggest that in such cases universal love allows, if not requires, acting unlovingly towards the one in order to act lovingly towards the other million. Nonetheless, love for the one would constrain one's action such that feelings of remorse as well acts of rectification would be warranted.

To sum up: An agent who is *never* willing to harm some in order to benefit others simply does not love everyone, at least not virtuously. An agent who loves everyone has a *pro tanto* desire to avoid harming anyone. However, when tragic circumstances arise such that the many can be saved from harm only harming some others, an agent who loves everyone would have a conclusive (all things considered) desire to harm some in order to save the many from a relevant harm.

7 Conclusion

In this dissertation I investigated the plausibility of explaining morality in terms of universal love. Minimally, a plausible account of morality will provide an adequate explanation of the moral properties of actions as well as the moral character of agents. The central question of this dissertation was: Can universal love play a foundational role in a plausible account of morality? I have answered this question in the affirmative. In particular, I have argued that universal love is well suited to play a central role in a virtue ethical approach to moral theory.

I want to conclude this dissertation by first summarizing my argument. I then discuss the limitations of the argument and some of the future work that remains to be done to develop and defend the moral theory sketched in this dissertation.

7.1 Summary

Briefly, I started this dissertation by sketching a volitional account of love (chapter 2). Using the volitional account of love, I developed an account of universal love (chapter 3). Further, I argued that universal love cannot play a foundational role in a deontological moral theory, but it can play a foundational role in a virtue ethical moral theory (chapter 4). I then developed a virtue ethical moral theory in which universal love plays a central role of *unifying* the moral virtues and is therefore conclusive for determining moral reasons for action (chapter 5). Lastly, I defended this moral theory by addressing some objections (chapter 6).

In chapter 2, I started by sketching an account of love for someone. On my view, to love someone is to desire the promotion of their well-being for their own sake. I understand desires as *pro tanto* dispositions to act: an agent *A* desires *x* if and only if *A*

has a *pro tanto* disposition to exercise her agency to bring about *x*. To love someone, then, is to have the *pro tanto* disposition to promote their well-being through the exercise of one's agency. This includes the *pro tanto* dispositions to (i) perform actions that one believes will improve their well-being, and (ii) refrain from performing actions that one believes will reduce their well-being.

Further, as a working assumption, I endorsed the so-called objective-list theory of well-being. That is, on my view, the elements of well-being are objective and plural — there are multiple elements that contribute to the well-being of any being of a given kind. I did not endorse any particular *exhaustive* list of the elements of well-being, though I did endorse some relatively uncontroversial elements of well-being. For instance, I have assumed throughout the discussion that pleasure, autonomy, knowledge, achievement, and friendship are (positive) elements of well-being for human persons.

In Chapter 3, I developed an account of universal love using my account of love. On my account, universal love is the desire for the promotion of the well-being of all beings with moral standing (that one veridically believes to exist). That is, universal love is the species of love that has as its object all beings with moral standing (that one veridically believes to exist). To manifest universal love, then, is to have the *pro tanto* dispositions to (i) perform actions that one believes will improve the well-being of all beings with moral standing (that one veridically believes to exist), and (ii) refrain from performing actions that one believes will reduce their well-being.

On my account of universal love, one need not be *acquainted* with a being with moral standing in order to love them. One could, instead, have a *de dicto* attitude of love for someone that one is not acquainted with as long as one believes that they exist (qua a

certain kind of being). Further, on my view, universal love is compatible with having a stronger desire for the promotion of well-being of one's close ones than the desire for the promotion of the well-being of strangers. Universal love is the desire for the promotion of the well-being of all beings with moral standing (that one veridically believes to exist), but this desire may be stronger for some beings than others.

Having developed an account of universal love in chapters 2 and 3, I addressed the central issue of this dissertation — the role that universal love can play in moral theory — in chapters 4 and 5.

In chapter 4, I argued that universal love can play a foundational role in a virtue ethical moral theory, but not in a deontological moral theory. I argued that universal love cannot be the object of a moral requirement because there cannot be a *duty* to love. This, in turn, is because love is not an attitude that is within an agent's voluntary control. As a result, universal love cannot play a foundational role in a deontological moral theory. I then suggested that the disposition to love universally is a moral virtue. Briefly, the disposition to love all beings with moral standing is a morally excellent disposition.

I used a target-centered account of the virtues to develop an account of universal love as a moral virtue. Central to the target-centered account of virtue are the notions of the *field* and the *target* of a virtue. The *field* of a virtue consists of the elements that are the sphere of concern of the virtue. A *virtue* for a given normative domain is a disposition to respond to elements within its field(s) in an excellent way, relative to that domain. The *target* of a virtue is what the virtue is aimed at within its field. The field of universal love, on my account, is the well-being of all beings with moral standing who stand to be affected, in one way or another, by one's agency. Further, the target of universal love, on

my account, is the promotion of the sum of the *weighted, relevant* interests of all beings with moral standing that stand to be affected by one's agency.

On my account, universal love justifies interpersonal aggregation of *relevant* interests, but not that of irrelevant interests. My account is thus a partial-aggregative view, rather than an anti-aggregative or full-aggregative view. On my view, in any given circumstance, certain *relevance factors* determine both (i) which interests are *relevant* for interpersonal aggregation, and (ii) the *weight* of the relevant interests. I don't endorse any specific position on the relevance factors, though some plausible relevant factors are the current level of well-being of the individuals, doing harm vs. allowing harm, whether the agent has a special relationship with any of the individuals involved, etc.

In chapter 5, I sketched a moral theory in which the disposition to love universally is the sole cardinal moral virtue that grounds all the other moral virtues. A cardinal virtue is a virtue that grounds some virtues in its domain but is not itself grounded by any virtue in its domain. I argued that the disposition to love universally is the sole cardinal moral virtue that grounds all the other moral virtues. In particular, I considered six commonly recognized moral virtues — benevolence, compassion, courage, honesty, respect, and justice — and argued that the field of universal love is wide enough to include the elements that these virtues are concerned. Thus, all of these commonly recognized moral virtues may plausibly be understood as different aspects of the disposition to love universally. In other words, the disposition to love universally may plausibly be understood as the grounds of all the other moral virtues.

Further, the moral virtue of universal love also unifies all the other moral virtues and is thus conclusively relevant for determining moral reasons for action. The

disposition to love universally, by virtue of being the sole cardinal moral virtue, also unifies the moral virtues by providing an overarching aim that helps adjudicate conflicts between the (non-cardinal) moral virtues. Accordingly, I sketched an account of moral reasons for action on which the disposition to love universally is *conclusively* relevant for determining moral reasons for action because it adjudicates conflicts between the non-cardinal moral virtues that ground *pro tanto* moral reasons for action.

Lastly, in chapter 6, I defended the central claim of this dissertation that universal is the sole foundational moral virtue that grounds all the (other) moral virtues but is not itself grounded in any other moral virtue.

7.2 Limitations and Future Work

Given the scope of this dissertation, there are some important limitations to its general argument. I now want to conclude by briefly discussing these limitations, in order to highlight the future work that needs to be done to develop and defend the moral theory sketched in this dissertation.

First, the plausibility of the claim that the disposition to love universally is the sole cardinal moral virtue depends, in part, on the plausibility of my account of love. I endorsed a particular account of love as the desire for the well-being of the beloved. Not surprisingly, there are competing accounts of love on offer in the literature. While I did provide some motivation for this view, a thorough defense of this account of love was beyond the scope of this dissertation. Of course, it is possible that my view of love is implausible, but the resulting moral theory is nonetheless plausible. Though, in that case I would have defended a rather different thesis than the one I have intended to defend.

Second, the plausibility of the moral theory sketched in this dissertation depends crucially on the plausibility of the account of well-being I have been working with. A crucial part of my argument for the centrality of universal love in virtue ethical moral theory was my claim that the moral virtue of universal love grounds moral virtues like respect and honesty. However, the moral virtue of universal love can ground moral virtues like respect and honesty only if the field of universal love includes elements like autonomy and knowledge. This, in turn, is possible only if autonomy and knowledge are among the factors that contribute to well-being. For, on my account, love is the desire for the promotion of the well-being of the beloved for their own sake. Love, then, can be concerned with autonomy and knowledge only if these elements are constitutive of the beloved's well-being. The plausibility of my overall argument, then, depends on the plausibility of the so-called objective-list theory of well-being that I have assumed throughout this dissertation. That is, I have assumed that there are objective elements, like autonomy and knowledge, which contribute to one's well-being, independent of one's desire for those elements. If it turns out that the most plausible theory of well-being is a hedonistic or a desire-satisfaction theory, then universal love could not ground moral virtues like respect and honesty since autonomy and respect would be outside its sphere of concern. A complete defense of the moral theory sketched in this dissertation, then, will also require a defense of the objective-list theory of well-being.

Third, the moral theory sketched here needs to be developed by providing a detailed account of the factors that determine the relevance and weight of competing interests for interpersonal aggregation. Recall that the target of universal love, on my account, is the promotion of the sum of the weighted, relevant interests of all beings with

moral standing that stand to be affected by one's agency. The relevance factors, whatever they are, collectively determine both (i) which interests are *relevant,* and (ii) the *weight* of the relevant interests, i.e., how much each relevant interest matters morally. In this dissertation, I merely suggested some plausible relevance factors like the current level of well-being of the individuals and doing vs. allowing harm. A fully worked out view would tell us (i) what the relevant factors are, (ii) how much each factor matters, and (iii) provide justification for (i) and (ii). Providing such an account is beyond the scope of this dissertation, though it is no doubt crucial for filling in the details of the content of the moral theory sketched here.

Fourth, the moral theory sketched here needs to be developed by providing a comprehensive account of the virtues. I have focused on the moral virtues, but a comprehensive account of the virtues would also discuss the prudential virtues as well as the intellectual virtues. I have understood the virtues as excellent dispositions, relative to a given normative domain. A comprehensive account of the practical virtues would address, not just the morally excellent dispositions, but also, the prudentially excellent dispositions. Importantly, the relationship between the moral and prudential virtues needs to be spelled out. Plausibly, there is general harmony but occasional conflict between the two. Virtue ethicists generally aspire to provide a comprehensive account of the practical domain, and it would certainly be desirable to do so. Relatedly, virtue ethicists across different traditions have recognized the importance of intellectual virtues like practical wisdom for acting and living well. Further investigation is therefore warranted into the role of the intellectual virtues for cultivating and manifesting universal love.

Lastly, the plausibility of the central claim of this dissertation — that universal love can play a foundational role moral theory — depends on the viability of virtue ethics as an approach to moral theory. For, I have argued that universal love is not well suited to play a central role in a deontological or consequentialist approach to moral theory. Instead, I have argued that universal love is well suited to play a central role in a virtue ethical approach to moral theory. On a virtue ethical approach to moral theory, all moral properties — goodness of character, states, and actions; and permissibility of actions — are grounded in the moral virtues. However, there are challenges to a virtue ethical approach to moral theory. If it turns out that virtue ethics is not a viable approach to moral theory, the moral theory sketched here would be implausible as well. Indeed, if it turns out that virtue ethics is not a viable approach to moral theory, we will have strong reason to conclude that universal love is not well suited to play a central role in moral theory at all. An important upshot of the general argument of this dissertation is that the issue of whether universal love can play a central role in moral theory depends on the viability of virtue ethics as such. To fully defend the claim that universal love can play a central role in moral theory, then, one would need to defend the viability of virtue ethics as such.

Bibliography

Alston, W. P. (1989). *Epistemic justification: Essays in the theory of knowledge*. London: Cornell Press

Badhwar, N. K. (1996). The limited unity of virtue. *Nous*, *30*(3), 306-329.

Badhwar, N.K. (2003). "Love". In LaFollette, H. (Ed.), *The Oxford handbook of practical ethics* (pp. 42-69). Oxford University Press.

Baron, M. (1987). Kantian ethics and supererogation. *The Journal of Philosophy*, *84*(5), 237-262.

Baron, M. (1998). Love and respect in the "Doctrine of Virtue". *The Southern Journal of Philosophy*, 36(5), 29-44.

Bloomfield, P. (2022). Virtues are excellences. *Ratio*, *35*(1), 49-60.

Brewer, T. (2018). Acknowledging Others. *Oxford Studies in Normative Ethics*, *8*, 9-31.

Brink, D. O. (1986). Utilitarian morality and the personal point of view. *The Journal of Philosophy*, 83(8), 417-438.

Brownlee, K. (2010). Moral aspirations and ideals. *Utilitas*, 22(3), 241-257.

Coady, C. A. J. (2008). *Messy morality: The challenge of politics*. OUP Oxford.

Cottingham, J. (1983). Ethics and impartiality. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, *43*(1), 83-99.

Cottingham, J. (1986). Partiality, favouritism and morality. *The Philosophical Quarterly (1950-)*, *36*(144), 357-373.

Crisp, Roger, "Well-Being", *The Stanford Encyclopedia of Philosophy* (Winter 2021 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/win2021/entries/well-being/>.

Dancy, Jonathan (1993). *Moral reasons*. Cambridge, Mass.: Blackwell.

Darwall, S. (2009). *The second-person standpoint: Morality, respect, and accountability*. Harvard University Press.

Dougherty, T. (2013). Rational Numbers: A Non-Consequentialist Explanation Of Why You Should Save The Many And Not The Few. *The Philosophical Quarterly*, *63*(252), 413-427.

Dorsey, D. (2012). Weak Anti-Rationalism and the Demands of Morality. *Nous*, *46*(1), 1-23.

Dorsey, D. (2016). *The limits of moral authority*. Oxford University Press.

Driver, Julia, "Moral Theory", *The Stanford Encyclopedia of Philosophy* (Fall 2022 Edition), Edward N. Zalta & Uri Nodelman (eds.), URL = <https://plato.stanford.edu/archives/fall2022/entries/moral-theory/>.

Frankfurt, H. (1999). Autonomy, Necessity, and Love, in *Necessity, Volition, and Love*, Cambridge: Cambridge University Press, 129–41.

Frankfurt, H. (1999). On Caring, in *Necessity, Volition, and Love*, Cambridge: Cambridge University Press, 155–180.

Frankfurt, H. (1999). *Necessity, Volition, and Love*. Cambridge University Press.

Hamlyn, D. W. (1978). The phenomena of love and hate. *Philosophy*, *53*(203), 5-20.

Helm, B. W. (2009). Love, identification, and the emotions. *American Philosophical Quarterly*, *46*(1), 39-59.

Helm, Bennett, "Love", *The Stanford Encyclopedia of Philosophy* (Fall 2021 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/fall2021/entries/love/>.

Hooker, B. (2015). The elements of well-being. *Journal of Practical Ethics*, *3*(1).

Hurd, H. M. (1998). Duties beyond the call of duty. *JRE*, *6*, 3.

Hursthouse, R. and Glen, P. Virtue Ethics. *The Stanford Encyclopedia of Philosophy* (Fall 2023 Edition). Edward N. Zalta & Uri Nodelman (eds.). URL = <https://plato.stanford.edu/archives/fall2023/entries/ethics-virtue/>.

Johnson, C. M. (2020). How deontologists can be moderate (and Why They Should Be). *The Journal of Value Inquiry*, *54*(2), 227-243.

Johnson, R. (1998). Love in vain. *The Southern Journal of Philosophy*, *36*(5), 45-50.

Kagan, S. (1992). The structure of normative ethics. *Philosophical perspectives*, *6*, 223-242.

Kagan, S. (2014). An introduction to ill-being. *Oxford studies in normative ethics*, *4*.

Kagan, S. (2019). *How to count animals, more or less*. Oxford University Press.

Kant, I., (1996). *The Metaphysics of Morals*. Trans. Mary Gregor. Cambridge: Cambridge University Press.

Kolodny, N. (2003). Love as valuing a relationship. *The Philosophical Review*, *112*(2), 135-189.

Liao, S. M. (2006). The idea of a duty to love. *J. Value Inquiry*, *40*, 1.

Liao, S. M. (2015). *The right to be loved*. Oxford University Press.

Penner, T. (1973). The unity of virtue. *The Philosophical Review*, *82*(1), 35-68.

Portmore, D. W. (2011). *Commonsense consequentialism: Wherein morality meets rationality* (Vol. 2). Oxford University Press.

Russell, D. C. (2009). *Practical intelligence and the virtues*. Oxford University Press.

Schroeder, Tim, "Desire", *The Stanford Encyclopedia of Philosophy* (Summer 2020 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/sum2020/entries/desire/>.

Setiya, K. (2014). Love and the Value of a Life. *Philosophical Review*, *123*(3), 251-280.

Setiya, K. (2022). Other People. *Rethinking the value of humanity*. Oxford University Press.

Silverman, E. J. (2019). *The supremacy of love: An agape-centered vision of Aristotelian virtue ethics*. Rowman & Littlefield.

Slote, M. (2001). *Morals from motives*. Oxford University Press.

Smith, N. R. (2018). Right action as virtuous action. *Australasian Journal of Philosophy*, *96*(2), 241-254.

Solheim, B. P. (1999). The Possibility of a Duty to Love. *Journal of Social Philosophy*, *30*(1).

Swanton, C. (2003). *Virtue ethics: A pluralistic view*. Oxford University Press.

Taurek, J. M. (1977). Should the numbers count?. *Philosophy & Public Affairs*, 293-316.

Toner, C. (2014). The full unity of the virtues. *The Journal of Ethics*, *18*, 207-227.

Trianosky, G. W. (1986). Supererogation, wrongdoing, and vice: on the autonomy of the ethics of virtue.

Um, S. (2024). Duty, Virtue, and Filial Love. *Philosophy*, *99*(1), 53-71.

Wallace, R. J. (2012, June). I—Duties of Love. In *Aristotelian Society Supplementary Volume* (Vol. 86, No. 1, pp. 175-198). Oxford, UK: Blackwell Publishing Ltd.

White, P. Q. (2025). Love first. *Philosophy and Phenomenological Research*, 110(3), 854-886.

Wolf, S. (1982). Moral saints. *The Journal of Philosophy*, *79*(8), 419-439.

Velleman, J. D. (1999). Love as a moral emotion. *Ethics*, *109*(2), 338-374.

Voorhoeve, A. (2014). How should we aggregate competing claims?. *Ethics*, *125*(1), 64-87.

Zhang, E. (2024). Individualist Theories and Interpersonal Aggregation. *Ethics*, *134*(4), 479-511.

VITA

Alok Tiwari was born in Dharupur, India on March 2, 1993. He received a B.E. in Electrical and Electronics Engineering from Sri Jayachamarajendra College of Engineering, Mysore in 2014. Shortly afterwards, he chose to pursue philosophy and received an M.A. in Philosophy from the University of Hyderabad in 2018. He then moved to the U.S.A. to continue his graduate career in Philosophy. He received an M.A and a Ph.D. in Philosophy from the University of Missouri in 2023 and 2025, respectively. He lives in Cedarburg, Wisconsin.