The University
Of
Sheffield.

The Practical Significance of the Second-Person Relation

By:

James H. P. Lewis

A thesis submitted in partial fulfilment of the requirements for the degree of
Doctor of Philosophy

The University of Sheffield

Faculty of Arts and Humanities

Department of Philosophy

Submission Date: 3rd January 2019

Abstract

Second-person relations are relations between individuals knowingly engaged in interaction with one another. These are the social contexts within which it is appropriate for one to think of and address another as 'you'. This dissertation explores the practical consequences for agents of relating to others in this fashion. A critical analysis is offered of Stephen Darwall's theory of moral obligations in terms of demands that can be addressed from the perspective of a second-person. On the basis of the criticisms raised, a broader conception of 'second-personal reasons' is advanced according to which there are a variety of species of practical reasons that are essentially grounded in second-person relations between agents, besides moral obligations. A paradigm case of such a species is the reasons that are presented in requests, and one chapter of this work is devoted to explaining this power people often grant to others: to intentionally create new, discretionary reasons for them. Drawing from several historical antecedents – particularly Martin Buber, Simone Weil and Emmanuel Levinas – the analysis of second-personal reasons is extended to include a discussion of the proper object of agapic love, and a discussion of the possible significance that face-to-face encounters may have for moral epistemology.

Acknowledgements

Contents

Chapter 1
Introduction

❖   You're in an airport, waiting to pick up someone you've never met before, holding their name on a sign, watching the faces of each new arrival as they glance at your sign. Each one moves straight on without paying you any further notice. But then someone comes through the Arrivals doors, sees your sign and looks up at your face.

❖   You're on the phone to a customer service line which, like other customer service systems you've called, plays you a recorded question, and computes your verbal answer using voice-recognition software before playing the appropriate follow-up question in the series. But then you realise that this last question was not a recording at all, it was the live, human voice of a customer service staff-member.

❖   You're running a marathon in a stream of countless other runners. The street is lined with crowds cheering you all on. But then in the midst of the crowd you hear the voice of a friend, exhorting you in particular, calling your name.

❖   You have crushed your bitter rivals at the Victoria sponge contest in the Baking and Vegetables tent of the annual farm show. You are looking forward very much to the reward of a long afternoon in the Refreshments tent of said show. But then your daughter tugs on your sleeve and reminds you that it is time for you to take her home.

The second-person relation is a way that people relate to one another. It is the relation within which people can appropriately address one another using second-person pronouns: *you, thou, thee*. It is a stance that people take towards their partners in distinctly interpersonal forms of interaction, like conversation. It contrasts with third-personal relations, where a person (or, as one might otherwise refer to the first-person, a *subject*) relates to objects, or relates to other people but as though they are mere objects. This contrast between the third-person relation and second-person relation is loosely illustrated by the suite of examples above. Each begins with relations that are in some sense third-personal, hinges around the phrase 'but then', and presents a contrasting second-person relation.

     This dissertation is about the *practical significance* of the second-person relation. That is to say, it is about the consequences that ensue from relating to other people second-personally. My focus is specifically on the way that certain *reasons* for action depend on the way that we intentionally relate to others (where this might involve our understanding of the way that they relate to us too). The fundamental idea is that some facts count as reasons for action only because of this special way that we relate to other people – second-personally – which contrasts starkly with the way that we relate to

other things that we encounter. Reasons like those that arise in promises, commands, demands, requests and favours, for example, are candidates for analysis.

In analysing the way that such reasons might depend on specific kinds of interpersonal relations, I aim to contribute to two issues of theoretical interest. One is in the field of ethics. It is the matter of whether (and how) our actions can be justified, in a way that delivers on our common conception of ourselves as rational beings. When we act on reasons that depend on relating second-personally to others, can we explain how this could be the *right* thing to do? Can the fact that we relate to others second-personally play a role in justifying our actions? (To pre-empt what is to come, I will offer some arguments supporting some positive answers to these questions, but other arguments against certain other kinds of positive answer.)

The other theoretical enterprise that this work contributes to is in the field of social philosophy. Understanding exactly how we relate to one another – how we think about, matter to, and influence one another – is a component part of understanding who and what we are. In philosophy, such a topic has sometimes gone under the heading of *intersubjectivity*. Part of my goal, then, in discussing the practical significance of the second-person relation, is to take a step towards understanding better what human intersubjectivity might consist in.

For now, in this introductory chapter, there are three goals, corresponding to the three sections below. The first will survey the somewhat disparate spread of ways that the second-person relation has been understood in recent philosophical literature. As a product of this survey, I will establish an understanding of that concept that will be drawn upon in the rest of the thesis. In the second section, I will motivate the project of analysing the practical significance of distinctly second-personal relations, and at the same time animate the general perspective on the topic that I will advance later. This will be done through the presentation of three vignettes – or mini-case-studies – which illustrate the explanatory gap in prevalent ethical theories that do not invoke the second-person relation. The final section, then, will summarise the five chapters that follow, offering a narrative in which, taken together, these chapters put forward a distinctive theoretical outlook.

1. What the second-person relation is

A body of philosophical and psychological literature about the second-person has emerged in the last twenty or so years. And on the face of it, this literature has the problem that not everyone seems to be talking about the same thing. There is ambiguity between authors over the necessary and sufficient conditions for a relation to be a second-personal one. Specifically, it is unclear whether a second-personal relation between two beings, A and B, requires:

  i)   that A and B are both actual people (rather than one being imaginary, or one being something other than a person);
  ii)   that they are engaged in interaction;

iii)   that they are engaged in cooperative joint action;

iv)   that they are emotionally engaged with one another;

and,

v)   that the attitudes comprising a second-personal stance (whichever they might be) are held by both, reciprocally.

In this section, I will set out and motivate some of the dominant views of what the second-person relation is, and spell out what these views have in common.[1] This is the shape of my survey. I will begin by explaining how the concept is thought of by several writers in the philosophy of language, before seeking to show how their conception is more similar than it might seem to that discussed in the philosophy of mind. Moreover, this same thread of commonality extends, I will argue, to account for *one* part of how the notion of the second person is discussed in ethics. At this point though, the consensus begins to falter, as researchers in psychology, and several philosophers of emotion, think of second-person relations as involving significantly more robust engagement between the A and the B than is assumed elsewhere. Indeed, this thicker conception is also at play, I suggest, in the *other* part of how the concept of second-personality is invoked in ethics. I will conclude this section with some remarks about how the variety of phenomena at issue can be aggregated in discussing the practical consequences of second-person relations.

A starting point for getting to grips with the philosophical discourse on the second-person is the role that interpersonal interactions are thought to play in the functioning of language. In the aftermath of Wittgenstein's *Philosophical Investigations*, attention was drawn to the nature of the social constitution of language. Wittgenstein set in motion a train of thought exploring the idea that a language is *comprised* of a set of rules that both govern, and are determined by, interactions between individuals who share some common endeavour, some form of life. In the context of this idea, the second-person relation might be thought of as that relation between speakers and interpreters that gives rise to social norms – paradigmatically the rules of a language.

Indeed, this conception of the second-person is the subject matter of Donald Davidson's 1992 article entitled 'The Second Person'. There, Davidson (1992, p. 269) makes explicit a version of that Wittgensteinian thought:

> [T]here would be no saying what a speaker was talking or thinking about, no basis for claiming he could locate objects in an objective space and time, without interaction with a second person.

What does Davidson have in mind when he refers to the second person? Well, his project in that paper is to address the social composition of a language – its basic interpersonal building-blocks.

---

[1] Conant and Rödl (2014, p. 4) raise a doubt about whether there is any fundamental commonality between what is meant by the 'second person' in philosophy of mind and what is meant by the same term in ethics. Whilst I sympathise with the complaint that there is some talking at cross purposes in the field, my survey here can be seen as an attempt to assuage their worry.

Accordingly, his interest in the relation between an I and a you is as the fundamental relation between speakers of a language. Of that relation, he says this:

> Each of them [the first- and second-person] must speak to the other and be understood by the other. They do not […] have to mean the same thing by the same words, but they must each be an interpreter of the other. (1992, p. 264)

As such, by referring to a second-person, Davidson means to refer to the perspective of one who can both *interpret* the utterances of the first-person (the speaker), and who can *be recognised* by the first-person as such an interpreter. It will be helpful for understanding later developments to get a feel for the motive behind each of these conditions.

The first of Davidson's conditions is that B occupies the status of a second person with respect to A only if B can interpret A's expressions. Indeed, the standpoint of an interpreter – of one who renders intelligible a subject's words and deeds, who therefore shares something of the subject's basic comprehension of the world – is central to all conceptions of the second-personal. This coheres with very basic intuitions about the kind of relation that anyone stands in to the being to whom they say 'you', or the being with regards to whom they *think* 'you'. Such utterances and thoughts evidently presuppose a relation not only of understanding, but of *mutual* understanding.

This is the second condition. Only if the first-person recognises someone else's capacity to interpret linguistic expressions, could that first-person ever form intentions – to speak in such a way as to be understood – which are necessary for linguistic transactions. So the second-person's capacity to understand the subject must be recognisable – and indeed, recognised – by the first-person in order for a second-person relation to be realised between them. The performances, interpretation of which are characteristic of second-personal relations, could only ever be produced by a subject who already takes her addressee *as* a competent recipient of those performances. Underlying Davidson's analysis of the rudiments of linguistic practices, then, is a notion of the second-person as recognised interpreter.[2]

A similar, parallel conception of the second-person, also born out of Wittgensteinian considerations about language and sociality, was advanced in Robert Brandom's 1994 work *Making It Explicit*. Whereas Davidson's paper sought to reconstruct the most minimal necessary social conditions for linguistic practice, Brandom's theoretical ambitions for the second-person are somewhat loftier.[3] Brandom advances the view that the meaningfulness of all linguistic expression – including the veracity of any given assertion – can be understood *just* in terms of the relationship between first- and second-person. The label in Brandom's theory for this second-personal approach is '*I-Thou* sociality'. Now it is beyond the present survey to fully explain – let alone assess – Brandom's inferentialist semantic theory. But what is very much relevant to an understanding of the development of the concept of the second-person is Brandom's idea that beyond being merely an interpreter of the

---

[2] See Bohman (2000) for a similar conception of the second-person as interpreter.
[3] Brandom (1994, p. 659 n.50) acknowledges Davidson as a 'notable exception', in being someone who appreciates the second-person relation as fundamental to linguistic practice.

subject's expressions, the Thou in *I-Thou* sociality must themselves also play the active role of a speaker.

To put the point simplistically, it is that the objective truth about any assertion – or, better, the very *idea* of such an objective truth – is arrived at by triangulating between the perspectives of pairs of interlocutors, each of whom have their own set of beliefs (Brandom, 1994, pp. 598-607). For example, if person A asserts something that person B takes to be false, then B acknowledges A's status as committed to this falsehood. The possibility of contrasting the perspective of any subject's beliefs with an objective perspective – from which the fact of the matter can be seen – is presupposed by B when she takes A to be wrong about something. And, of course, B can only take A to be wrong about a claim that A expresses if B herself also has some view on the matter. Second-persons, for Brandom, do not merely understand what first-persons say, they *evaluate* their claims. And any such evaluation can only be made from the point of view of an *active* player in what Brandom (Brandom 1994, p. xviii) famously calls 'the game of giving and asking for reasons'.

Davidson's model of the second-person was as interpreter, Brandom's is as interlocutor. This development has been influential. A school of enquiry has grown in the ensuing twenty years elaborating on the philosophical importance of the kind of second-personal interaction that Brandom locates at the bedrock of discursive practice. For instance, Glenda Satne (2014, 2017), Jeremy Wanderer (2010), and, in their collaborative work, Rebecca Kukla and Mark Lance (2009), might all be thought of as exemplars of this school. They all take up the idea that there is some issue of philosophical significance to be analysed in the relations between pairs of individuals who both understand and evaluate one another's performances in accordance with the norms of a social practice. (Incidentally, a particular language – like English – is the sort of thing that is meant here by a 'social practice', but similar second-personal interactions of holding one another to the norms of the game might also be thought to characterise various non-linguistic practices: actual games, like chess; cooperative activities that are structured by tradition and expectation, like, say, carpentry; and indeed other interactive but distinctively interpersonal activities, like ensemble music.)

So on the one hand, the way that Brandom understands the second-person relation has been important in the development of this literature partly because a number of subsequent scholars have shared his basic understanding, and have elaborated on it. But on the other hand, Brandom's innovations have another kind of significance: by causing a fault-line to emerge that foreshadowed a divide between two deeply contrasting ways of conceiving of the second-person relation. Indeed, as will become clear in the next section of this chapter, that divide creates the space that much of this dissertation aims to occupy.

The fault-line in question was opened up in an exchange of articles between Brandom (2000) and Jürgen Habermas (2000). Habermas alleges that the relation between a speaker and someone who understands and evaluates their utterances is not truly a second-personal one (Habermas, 2000, p. 345). The basis for this complaint is that such a relation does not depend, sufficiently, on *inter*action. It is a

requirement of Brandom's *I-Thou* sociality that the second-person who understands the first-person must themselves be an active participant in the practice, with an evaluative perspective built out of their own commitments to certain beliefs. However, it is not a requirement of this picture that any given interpreter, in order to count as a *Thou* in the relevant sense, must actually be disposed to react to the particular speaker whose utterances they interpret and evaluate. That is, it is sufficient for a relation to be an *I-Thou* relation in Brandom's sense that the second-person would be *capable* of reacting to the first-person in interaction in the manner of their social practice. But they do not need *actually* to react. The action does not have to be symmetrical. Habermas' criticism, therefore, stems from the sentiment that properly second-personal relations necessarily involve some kind of face to face interaction, which he thinks typically has the goal of some sort of cooperation (Habermas 2000, p. 346).

There is thus something of a chasm in conceptions of the second-person relation between those (like Brandom) who have in mind the relation between a subject and any other person to whom they could address their words and deeds, and on the other hand those (like Habermas) who have in mind relations between a subject and a particular other person who, standing before the subject, offers their own words or deeds for evaluation in return. For the time being, it will be best to make a note of this division of opinions – which, as mentioned, will recur later – but for now to leave this to one side. The motive for leaving this disagreement to one side for a moment is to show that notwithstanding Habermas' critique, there is a notable, extensive degree of commonality between the conceptions of the second-person relation shared by a number of philosophers over the course of the development of these debates. Habermas aside, there is a coherent body of work in philosophy of language that began in the early 1990s which finds a central place for a notion of the second-person relation as that between a speaker and their addressee – one capable of evaluating their utterances. Furthermore, as will now be seen, such a notion of the second-person is broadly taken up by the contributors to the more recent debates about second-personality in the philosophy of mind.

One such debate about the second-person has been fought over whether expressions involving second-person pronouns are expressions *of* a distinctive *kind of thought*. In other words, this is a debate between those who think that all thoughts can be completely understood as thoughts of a subject relating themselves to various objects (including other people as objects), and on the other hand, those who do not. Thus Richard Heck (2000, p. 12) and Christopher Peacocke (2014) have claimed that thoughts that are represented in language with second-person pronouns – thoughts such as 'you are cool' – are in fact reducible to arrangements of third-personal thoughts – such as 'the person to whom I am speaking is cool'. But others, such as Sebastian Rödl (2007, chpt. 6; 2014), Jane Heal (2014), and Adrian Haddock (2014), advance the opposite view: that a thought like 'you are cool' is never reducible to third-personal thoughts. They think that such a thought expresses a mode of thought wherein others are represented not as mere objects but also as second-persons.

To proceed with this survey of the various conceptions of the second-person, I will now summarise aspects of Heal and Rödl's views. Being the most worked-out positions in this neighbourhood of the literature, they are illustrative of the strengths of the view that second-personal thought is distinctive and irreducible, as well as exemplifying how the second-person relation is understood in this domain.

The starting point for Heal in thinking about the second-person bears an important resemblance to the notion as construed by Brandom et al.

> When one person utters a remark including 'you', in the central case, he or she addresses another person, where 'addressing' is an open act, something the occurrence of which is common knowledge between the addresser and the addressee…. Such open addressing of another is central to the second personal character of what occurs (Heal, 2014, p. 320).

To say that 'open addressing' is central to the second-personal character of a transaction implies that for a relation to be second-personal *is* for it to be such that it carries 'open addresses'. And all Heal means by an address being 'open' is that it is known to be an address by both addressor and addressee. Which is to say that the addressee must be capable of understanding the addressor, and the latter must be aware of this. As such, one can readily see that what Heal has in mind as a second-person relation is akin to the aforementioned idea of relations of mutual interpretation.

But Heal goes on from this starting position to argue that distinctly second-personal thoughts could arise out of one subject and another cooperating. Her animating idea is encapsulated in the slogan ''you' is 'we' minus 'I'' (2014, p. 317). She advances the claim that what it is to relate to another second-personally, is to relate in an 'addressing' fashion to them (either as the addressor or the addressee: the relation is second-personal in both directions) in the context of some joint action. To take an example, then, suppose that Eira and Meinir are lying in ambush. They have agreed that when Eira sees the convoy arrive at the gorge, she will signal to Meinir to spring from her position, after which Eira herself will spring from her own position. In this context, Eira's signal would be a manifestation of a distinctly second-personal thought. The signal communicates a thought of the form, 'Now Meinir! You spring from your position'. Heal's point is that the 'you' in such a thought could not be substituted for any third-personal reference to Meinir. The essence of the signal would be lost if one were to replace it with 'Meinir should spring from her position'. In thinking of Meinir as 'you', in the context of their shared action, Eira really thinks something more like '*We* should spring from Meinir's position' – or so Heal argues. The second-person relation, then, which is manifested in second-personal thought, would be intrinsic to cooperative social interaction.

This same line of thought is, in a sense, taken further by Rödl (2014). One could say that it is the same line of thought that is taken up, in that Rödl agrees with Heal that second-person relations exist in the context of cooperation, and that their paradigm is the act of open address. But the sense in which this line of thought is taken further is that Rödl claims that all instances of address – which he calls 'intentional transactions' – are themselves already cooperative actions which always express irreducibly second-personal thoughts.

Rödl's argument for the claim that all instances of address express irreducibly second-personal thoughts is based on his view of the way that the addressee is represented by the addressor in an act of address. The addressee is represented in such a way that if, after their successful ambush, Meinir says to Eira, 'you're cool', then Eira would legitimately hear this as '*I* am cool (according to Meinir)'. That is to say, the second-person is represented in address in such a way that it is *equivalent* to the first-personal expression of the addressee. The crux of Rödl's analysis is this equivalence. Since A's representation of B as 'you' is equivalent to B's representation of themselves as 'I', it follows that when something is predicated of the other in address, this is a kind of *self*-predication, namely 'dyadic self-predication' (2014, p. 310).[4]

This is all to say, whether or not it ultimately turns out to be true that there is an irreducibly second-personal 'mode of thought', that the kinds of relations between people that such a mode of thought might arise from do exist – unquestionably. Those are second-personal relations which are paradigmatically represented by instances of intentional address, and which make certain kinds of joint action possible. They require two individuals who are capable of understanding one another to mutually acknowledge that mutual capacity. What is relevant for present purposes, is that the basis of the second-person relation as it is understood by proponents of the idea of second-personal thought is broadly consistent with how that relation is conceived in the earlier debates in the philosophy of language.

So far, in the spirit of colouring in the background to the key term in my title, I have been discussing areas of philosophy about which this dissertation will in fact have little more to say. I am not going to advance any views about the second-personal basis of linguistic normativity, nor about the irreducibility of second-personal thought. What I *am* going to attend to is ethics: the role that second-person relations play in shaping how we should act. That is a domain of philosophy that has also featured discussion of the second-person in the past couple of decades; indeed, this sub-discipline has been the locus of the most extensive such discussion. Debate has revolved around Stephen Darwall's work (2004, 2006, 2013a, 2013b), which introduced the focus on second-person relations to modern Anglophone moral philosophy. I will take up and assess Darwall's views in greater detail in the next chapter. For now, continuing the survey of the different things that are understood by the notion of the second-person relation, I shall offer a sketch of his perspective on that matter.

Two things are worth noting about Darwall's view of what the second-person relation is. The first is that it has been singularly influential in informing how the concept is conceived in practical philosophy: there is by now a large body of secondary literature spawned by his 2006 book *The Second-Person Standpoint* (hereafter *SPS*). The second is that the book defines the second-person relation only quite loosely, and thus admits multiple competing ideas of exactly what kind of interpersonal relation is at the heart of his claims. Two such competing views are particularly salient:

---

[4] The same point is made by Haddock (2014, p. 347) when he says that, 'because Y's recognition [of X's address] *does* represent Y in the first[-]personal manner, so does X's volition [to address Y]'.

the idea of the second-person as an imaginary perspective available in individual reflection, and the incompatible idea of the second-person as a particular concrete individual who is actually engaged in mutual recognition with the subject. I will explain the first of these views and show how it coheres with preceding moments in the history of this concept. Then I will contrast it with the second view, and illustrate some of the historical precedent for *that*. I will conclude my survey then with some reflections about the tensions within the concept of the second-person relation and the fuel that the concept may nonetheless provide for further thinking about ethics.

Like quite a number of the figures already mentioned, Darwall does not primarily focus on the second-person *relation*, but rather, in his case, on the second-person *standpoint*. In light of which, I should say, though, that any notion of a second-person standpoint, or of second-personal sociality, or second-person thought, all necessarily presuppose a *kind of relation* between a first- and second-person. So I am justified in re-assembling this disparate array of ideas under one roof in the concept of a second-person relation. Indeed, I am not the first to try to do so (see Eilan, 2014). Nonetheless, it is notable that whereas a relation is something that holds between two people, a standpoint is occupied by only one. As such there is some ambiguity built into the definition that Darwall gives at the outset of his book (Darwall, 2006, p. 3):

> Call the second-person standpoint the perspective you and I take up when we make and acknowledge claims on one another's conduct and will.

The ambiguity here is over whether either you or I can be said to have taken up the second-person standpoint if the other has not. That is, if I make a claim on your will, but you don't acknowledge it, am *I* then occupying the second-person standpoint without you? Darwall's answer is interesting. His view, I think, is that the second-person standpoint *implies* mutuality, which is to say that any claim that one might make or acknowledge from that standpoint is such that it calls for recognition by both addressor and addressee; this is *not* to say, however, that any such claim actually attains such mutual recognition. Nor, therefore, does Darwall's central conception of the second-person relation apply only to pairs of individuals who are actually engaged in any kind of relation of mutual recognition or understanding.

Rather, then, the second-person relation that Darwall has in mind is between a subject and the *stance* of a (real or imagined) possible addressor or addressee. By defining the second-person standpoint (and, thus, the second-person relation that the standpoint is within) in terms of *claims* that people make of one another, Darwall echoes something of Brandom's perspective. For both Darwall and Brandom, what is essential to the second-person relation is that the parties are not merely passively interpreting one another. They must each have an active stance on what the other does. For Brandom, this was manifested in an understanding of speakers holding one another to account for adherence to linguistic norms, and to the facts about which they purport to make claims. For Darwall the picture is much the same except that linguistic norms are replaced by moral ones. In order for our

thoughts and attitudes to be second-personal, they must imply more than a mention of 'you', they must imply a thought that begins 'you should…'.

Relating second-personally therefore means engaging with the other's perspective as an agent. From the second-person standpoint, the deeds of another are not *mere happenings* – external events, outside of one's control, which one can only watch as they determine the conditions of one's own agency (which is to say, as they determine what it is possible for one to do). Rather than being mere happenings of this sort, from the second-person standpoint, the deeds of another are *doings*: they are actions which must be justified with reference to reasons. Like mere happenings, doings are events, but unlike mere happenings, they are events that can be right or wrong. And whether they are right or wrong will have consequences for the agents responsible. On Darwall's picture, relating to someone second-personally is precisely the stance from which one treats them as responsible for the events that they intentionally bring about, that is, for their actions (Darwall, 2006, pp. 11-12).

By tying the concept of the second-person so closely together with the idea of holding agents responsible, Darwall also ties it to a prior motif in moral philosophy: namely, Peter Strawson's notion of the 'participant stance' (Strawson, 1968, p. 79; Darwall, 2006, p. 69). As well as being a relevant detail in the emerging historical backdrop to the development of this concept, the link with Strawson's participant stance casts some light on the composition of the second person-relation as conceived by Darwall. What is revealed in this connection is the symmetry between the attitude that one takes towards another's action, and the attitude that one takes towards one's own. If the second-person relation just is that from which one sees certain events as the doings of agents, and from which one thus treats such agents as responsible for such events, then one must take something very much like a second-personal stance towards one's own deeds. Indeed, it *is* very much alike, but not exactly alike. Relating second-personally to others means being disposed to feel certain feelings of approval and disapproval in response to their deeds, such as blame and resentment. Treating oneself as a responsible agent also comes along with dispositions for such reactive attitudes, but with the difference that the repertoire of attitudes is modulated into a self-reflexive key. Instead of feeling the anger of blame and resentment, second-personal attitudes towards one's own wrong-doing yield the anguish of shame, guilt and embarrassment.

So Darwall's conception of the second-person relation is such that it makes sense to talk of taking a second-person standpoint towards oneself (Darwall, 2006, p. 78). An exchange of articles with Christine Korsgaard (2007), drew attention to the fact that on Darwall's view, a second-personal relation can be constructed imaginatively. Korsgaard (2007, p. 11) claims that:

> because of the reflective structure of human consciousness […] every rational agent stands in what Darwall would call a second-personal relation to herself – she has a second-personal voice within.

And in his response, Darwall (2007, p. 55) acknowledges that whenever I as an agent am deliberating about what to do,

I have to keep track of things from my perspective on you but also simultaneously from your perspective on me.

This version of the second-person relation becomes altogether clearer when this acknowledgement in response to Korsgaard is coupled with Darwall's view that the 'you' whose perspective must be considered in all of one's deliberations is not any particular person or set of persons, but rather the perspective of *anyone at all*. In *SPS* the second-person standpoint need not be the perspective of a particular person, but is the standpoint of anyone at all with whom one could in principle take up mutual participant stances. Darwall (2006, p. 35, my emphasis) says:

[second-personal competence] must consist in something like the capacity to choose to do something only if it is consistent with demands one (*or anyone*) would make of *anyone* (hence that one would make of oneself) from a standpoint we can share as mutually accountable persons.

Putting these pieces together, then: a second-person relation as construed by Darwall in *SPS* is a relation between any two agents who relate to one another from the participant stance, where, thus, each treats the other as responsible for their conduct. The two poles of a second-person relation need not be actually engaged with one another in any way, indeed, they need not both be *actual* at all; at least one could be a merely *possible* person. Suppose that Mared is in the process of dumping her old, broken freezer in a country lay-by to avoid paying to leave it at the skip. Even though she is totally alone, as she lugs it out of the back of her van, she conjures an image of a stranger catching her in the act. Her imagination of the stranger is completely vague – in her mind's eye she does not see a person with a particular body, or a certain physiognomy; she does not imagine hearing a voice with any particular tonal qualities. But the thought of a stranger, any stranger, seeing her fly-tipping makes her feel horribly guilty. She realises that she shouldn't really be doing it (in a way that she had not thought of it before). This relation between Mared and her faceless, voiceless, fictional stranger, is in a sense a paradigmatic instance of a second-person relation on Darwall's view in *SPS*.

There is a certain extent to which the version of the concept in *SPS* chimes with other ways that the second-person relation is thought of in the literature. Specifically, to view the first- and second-person perspectives as really nothing more than two sides of the same coin, as sharing their fundamental structure, is harmonious with Rödl's view presented earlier. However, despite chiming with some aspects of how the second-person is thought of, the conception of *SPS* also stands in unequivocal conflict with other trends in the literature. Such conflict stems from two features.

One is relatively trivial: in order to advance his particular theory about moral normativity, the definition of second-person relations in *SPS* is restricted to the dimension of *moral* interpersonal attitudes. This runs contrary to how the relation is thought of by many if not all of the theorists mentioned above whose primary focus is on the philosophy of language, or the philosophy of mind. But this is a trivial difference because the restriction to moral matters turns out only to be a technicality introduced in *SPS*. In later work on the second-person, Darwall considers the second-personal qualities of certain non-moral interpersonal dynamics such as trust and love (Darwall, 2016, 2017, respectively). Evidently then, the restriction of second-personality to moral attitudes was only a

surface-level specification, masking the fact that the phenomena at issue really has both moral and non-moral dimensions.

The other conflict, though, between Darwall's conception of the second-person and a rival trend in the literature, is over whether second-person relations could be anything other than the product of real interaction between real people. The trend in question, with which Darwall's view comes into conflict here, is one that understands the second-person relation as that between two people when they interact, and which analyses the states of mind that (always) arise in such interactions. This trend has been comprised predominantly of work done by psychologists and philosophers of psychology, but is also representative of the interest in the second-person relation of those working on the phenomenological tradition. Its major point of difference with Darwall's view – to be clear – is that according to this trend of scholarship the phenomena of interest that goes under the label of the second-person relation is necessarily an *actual* interaction between people.

Since many of these researchers are psychologists and phenomenologists, their interest is not primarily in the normative consequences of relating to someone a certain way (second-personally), but in the mental activity associated with those relations. In this spirit, as early as 2001, Dan Zahavi (2001, p. 166) urged that an understanding of the nature of both subjectivity and intersubjectivity depended on some better understanding of second-personal interactions. And a body of work has developed that seeks to deliver on this call of Zahavi's, with a range of analyses of different aspects of the phenomena bound up in second-personal interaction. For example, Shaun Gallagher and Daniel Hutto (Gallagher & Hutto 2008), among others, have put forward an 'Interactionist Theory' of social cognition, based on observations about the development of intersubjective understanding through embodied interaction – which is to say, in the idiom of this tradition of enquiry, through second-person relations. Similarly, Ingar Brinck (2008) analyses the gaze of infants into the eyes of others as one among several kinds of pre-linguistic second-personal interactions.

The research trend in question is drawn together in a 2013 article calling for a shift in neuroscientific research, towards a paradigm in which mental states relating to social life – thoughts about other people, and so on – are studied through the lens of actual interactions between people, which the authors call a 'second-personal neuroscience' (Schilbach et al. 2013). The research agenda in this burgeoning field begins with the observation that,

> social cognition [that is, the formation of beliefs about other people] is fundamentally different when we are in interaction with others rather than merely observing them (2013, p. 393).

But the conception of second-person relations that such scholars have in mind is not just any relation between interacting individuals. A further necessary condition for such an interaction to constitute a second-person relation seems to be that the parties to the relation are in some sense 'emotionally engaged' with one another (2013, p. 396). This should not come as a surprise. What unites the psychologists, phenomenologists, and philosophers of emotion in their focus on second-person relations is an interest in the experiences that come alongside second-personal interactions. That is,

these research programmes set out to study what it *feels like* to relate to others second-personally, and indeed, what is going on in our bodies when it does feel like that.

So there is Darwall's conception of the second-person, which can be abstracted from actual social interactions, and there is a rival conception which cannot. Given their methodological differences, and the divergent explananda of their theoretical projects, it is tempting to think that the conflict between these two conceptions of the second-person is a totally insignificant terminological issue: that the ethicists are interested in one thing, the psychologists another, and the fact that they happen to be using the same term is an unimportant coincidence. It may well be, that is, that there is only a superficial overlap in the subject matter of the two enquiries – the one into the normative sentiments derived from treating one another as responsible agents, the other into the emotions and other mental states arising from social interaction.

But this is certainly not the whole story. There is a final complication in this narrative of the recent history of the concept of the second-person. This complication is introduced by a perspective within the discourse in moral philosophy which contends that the sorts of reasons that we are well-placed to address to one another when we relate to each other second-personally (the sorts of things that Darwall calls 'second-personal reasons'), are *inextricably* bound up with actual relations between real individuals. This perspective has been advanced most prominently by Michael Thompson (2004) and Jay Wallace (2007). In slogan form, this school's essence might be put thus: second-personal ethics depends on second-person phenomenology (although I do not suppose either Thompson or Wallace would put it quite that way.)

Thompson's influential essay introduces the notion of 'bipolar' relations, and the moral reasons that arise in them. Bipolar relations differ from Darwall's notion of second-person relations in that the former must be actual relations between two particular individuals. Darwall's notion of the second-person standpoint is such that the subject is at one pole of the relation, and at the other is an abstract notion of a second-person, a position that could just as well be occupied by any member of the moral community. By contrast, in bipolar relations, the subject relates to a particular other. The poles of bipolar relations are occupied, paradigmatically, by a promisor, and a promisee, a wronger and a wrongee: by someone who has a duty to do something, and by the one to whom they owe that duty. It is because bipolar relations, unlike Darwall's second-person relations, always involve two actual people, that one can understand the dynamics of bipolar relations (but not Darwall's second-person relations) in light of the emotional and other psychological phenomena that are always attendant in instances of actual interpersonal interaction. But Thompson's idea is also similar to Darwall's, in that he too is primarily concerned to spell out the ethics of these distinctive interpersonal relations. He too identifies a species of practical reason that is grounded in the relation that people stand in to one another: namely, Thompson explores the possibility that an individual's moral rights and the parallel duties that those rights imply for others, are a function of the bipolar relations that that person stands in in relation to those others. Whether moral duties ought to be understood as generated

by bipolar relations, rather than second-person relations in Darwall's sense, is the matter that Wallace takes up. Darwall has subsequently responded with an essay called 'Bipolar Obligation' (Darwall 2012) in which he distinguishes the claims originating from particular relations between particular individuals, from the moral claims that originate from the relations that hold between any morally responsible individuals whenever they take the second-person standpoint.

This introductory survey is not the place to delve any deeper into the debate there over the success of Darwall's defence against the challenge posed by Thompson and Wallace. But what is relevant for present concerns is that even though the kind of social connection that Thompson draws attention to is given a different name, it is far closer to what the phenomenologists and psychologists refer to as the second-person relation than Darwall's own conception is. From my perspective, this is interesting.

This is because of an animating idea that runs through this dissertation: that the subject matter under investigation in all of these variegated developments about the second-person is actually more unified than it might seem. That is, informing the perspective of my discussions in the chapters below is the hunch that there is, at root, *something that it is like* to relate to one another second-personally. And more than this, a live option that I propose to explore is that the reasons that guide our actions in our social lives are often *grounded* in various ways *in* that something. That is very vague. I hope that a more detailed version of the same conviction will emerge as we go along.

I also hope, by now, that it is clear that it would not be possible to give a single definition of the second-person relation without precluding some of the range of views that I have laid out. So rather than attempting such a singular definition at the outset of this work, I will be satisfied to have presented this survey of the spectrum of conceptions. To recapitulate, there are a handful of different sub-disciplines in philosophy within which the 'second-person relation' has become something of a buzz-word, where there is hope that new insights on old questions can be gained by attending to that relation. And there is some commonality between these sub-disciplines in what is meant by the second-person relation. Everyone seems to converge in thinking, with Davidson, that two people relate second-personally when they recognise each other as *understanding* one another. This is perhaps the most fundamental characteristic that distinguishes second- from third-person relations. Most if not all figures surveyed further agree that within the context of such mutual understanding, the paradigm expression of a second-person relation is the act of address. This is some action intended by the first-person to communicate something to the second-person, and for the latter to recognise this very intention. Interpersonal address presupposes mutual understanding.

However, the uncontentious common ground does not extend much further. On this basis some – like Darwall and Brandom, but perhaps also Heal and Rödl – think of the second-person relation as a *structural* feature of the relation between one subject and another. In being such a purely structural feature, the second-person relation can be recreated in thought alone by imagining standing

in an interpersonal relation with that particular structure (namely, the structure of, say, understanding and evaluating the words and deeds of another who understands and evaluates your own in return).

But on the other hand, others construe the second-person relation as necessarily *material*, and concrete. This was already implied by Habermas, but developed further by Zahavi (who invokes support from various figures in the phenomenological canon), augmented with scientific content by the researchers behind the new second-personal neuroscience, and echoed again by Darwall's critics who insist on the 'bipolarity' of moral reasons. These diverse proponents of the material conception of the second-person are united by the thought that whatever it is that distinguishes second-person relations, it is inextricable from actual interactions between real individuals. The emotions and normative intuitions that accompany such real interactions are the true phenomena at issue, they say. Of the many more and less subtle differences in how the notion of the second-person relation has been deployed, then, this dichotomy between structural and material conceptions is both the deepest, and the most prevalent across several areas of enquiry.

In the following chapters, it will be important, of course, to be clear about which understanding of the second-person relation is meant whenever any is invoked. But in general I will not be restricting my focus to one conception or another. My goal is to develop a better understanding of the practical consequences of relating to others, as we often do, second-personally. To do this will sometimes require playing different conceptions of the second-person relation off against one another with a view to exposing their underlying basis. In the next section, now, I will offer some further explanation of the theoretical calling of this project, and how it will best be undertaken.

2. Explanatory gaps

I have offered an answer to the question of what the second-person relation is. Admittedly, it was a complicated answer, but it was not a simple question. Now I would like to motivate my endeavour of assessing that relation's practical significance. To do so, I will sketch three well-known topics in moral philosophy and set out to show that each conceals, just beneath the surface, a myriad of starkly under-theorised issues about the nature and consequences of our second-personal relations with one another. As well as showing why there is a need for this theorising, I hope to convey through these three examples a sense of the perspective on moral philosophy that will be characteristic of the approach of the coming chapters. The three topics that I propose briefly to discuss are promises, reactive attitudes, and moral obligations.

Promises are a natural beginning point for two reasons. On the one hand, the reason we have to keep our promises has been a topic of much philosophical discussion since at least the early-modern contract theorists. And on the other hand, any such reason quite evidently has everything to do with the relation between two mutually understanding people engaged in an act of address, so is, on anyone's account, a practical consequence of a second-person relation. A further notable feature of

the discussion about promises is that most everyone agrees that we *do* have reason to keep our promises – at least in some cases – which is to say that these second-person relations *actually* play a role in the creation of practical reasons.

A natural question is *when* promises create reasons. There is a range of theories answering this question. I don't wish to get too much into the intricacies of those theories, so instead it is possible to sketch something of a crude answer, to be taken as an example. Suppose, then, that A creates a reason for themselves to φ, by making a promise to B just when:

i)   A communicates to B that B can trust A to φ

ii)   A intends for (i) to count as a reason for herself to φ

iii)   A intends for B to recognise (ii)

iv)   A intends for B to recognise (iii)

v)   And B does recognise (ii) and (iii).

For my purposes, nothing hangs on whether this is a totally accurate account of promising, so long as at least the spirit of it is representative of how promises are often thought of by philosophers.[5] Specifically, I mean to illustrate the idea that whatever reason promises create depends on some kind of recognition, and possibly some higher-order recognition, between promiser and promisee.

Most philosophers – regardless of the moral theory that they subscribe to – believe that promises create reasons. And this belief is generally a straightforward product of philosophical intuition: the intuition that, other things being equal, it is wrong to break a promise. This intuition, plainly, is drawn from our experience of ordinary interpersonal life – our experience of making promises, the guilt of breaking them, the hurt and anger of having promises broken that were made to us.

So if on the basis of this sort of intuition one believes that promises create reasons, and one also subscribes to roughly the spirit of the above account of when a reason-creating promises is made, then one is committed to the idea that recognition of intentions can play a role in generating reasons – possibly that recognition of recognition of intentions can too. But the mere fact that you know what I intend to do surely cannot *by itself* create a reason for action. How could the matter of what I should do be at all determined by the apparently independent matter of what you happen to know about what I intend to do? Rather than independently favouring certain actions, the fact of another's recognition of one's intention plays a role in creating a reason only in conjunction with some matrix of other facts. Specifically, those are facts about B's subsequent *trust* in A φ-ing; B's *interest* that could be harmed if her trust is betrayed; B's *standing* as deserving fair treatment; A's standing in B's eyes in the event of A either betraying B's trust, or consequently harming her interest, or both; A's standing in her own eyes, and the eyes of other people, in the event of her wronging B; the value of the relationship between A and B; and so on.

---

[5] For a range of expressions of how promises *are* thought of by philosophers, see Sheinman (2011).

This matrix of factors is what generates reasons when promises are made successfully. But the question that I want to raise, which leads to the central point of this case-study, is this: what about when attempted promises are unsuccessful? Various moral philosophical theories will declare that these factors are alive, and interlinked, when promises succeed in generating reasons – and not just any old reasons, but moral obligations. Whereas the rest of the time, when promises are not made – or not made successfully – most moral theories have nothing at all to say about how this same matrix of factors in the relationship between mutually recognising individuals determines what either should do.

To illustrate the point, consider Tamsin and Beth, friends who have both been invited to a fundraising event that Tamsin is keen to go to – very much unlike Beth. Suppose first that Beth, knowing that Tamsin wants to go, and that Tamsin want her company too, promises Tamsin that she will go. In this case, notwithstanding any excuses that might arise, one would think that Beth has created for herself a reason – maybe even a moral obligation – to honour her promise and go to the fundraiser. Moreover, the grounds of this reason – the facts that *make* it a reason – are some arrangement of the following: Tamsin trusts Beth to go. Tamsin is entitled to expect Beth to go. Tamsin will be happy if Beth goes. But, of course, she will be disappointed and unhappy if she doesn't. Tamsin will be angry with Beth if she doesn't. If she doesn't go, Beth's very character will be that much less worthy of moral praise, that much more worthy of moral blame, from Tamsin's perspective, from Beth's own, and indeed from anyone else's. Tamsin may actually *like* Beth less if she lets her down. Indeed, perhaps Tamsin would be right to like Beth less – perhaps her failure to keep a promise to a friend would either constitute or indicate a reduction in Beth's love-worthiness, at least from Tamsin's perspective, but, again, maybe from anyone else's too.

One interesting thing about Beth being sensitive to these sorts of considerations, that undergird the reason she has to keep her promise, is that some of these factors may already have been at play earlier on, before she made her promise. But by making a commitment to Tamsin of which the latter is aware, this sensitivity to those considerations is significantly intensified. But now let's alter the case slightly. Suppose that when Beth resolves to make her promise, which she makes by phone, the phone signal is interrupted just at the moment in the call when she utters the promise. So, unbeknownst to Beth, Tamsin doesn't hear the promise at all. In this altered scenario, according to the rough account of promising proposed above, no promise has been made. As such, from the point of view of many moral theories, those factors that so nearly came into play in Beth's practical reasoning, remain inert. Or, if not inert, in the absence of a successful promise, there is at least no systematic way for ordinary moral theory to account for the interaction between those factors. There is no developed theory that says *how* those factors might be connected up, and thus how they might bear on what Beth ought to do.

This is an illustration of an explanatory gap that would be filled by a thorough account of the practical significance of the second-person relation. What such an account would explain, would be how such things as interpersonal judgements of others in moral terms have a bearing on interpersonal

judgements in non-moral terms (in terms, say, of loving and liking each other). It would explain the way that when we relate to another as a person, rather than as a mere thing, we occupy a practical perspective from which this matrix of considerations becomes both salient and ordered. Darwall's theory, which I will discuss in Chapter Two, is one such attempt to explain the second-personal normativity that this example brings to light. In Chapter Three I will further contextualise the second-personality of promises in relation to the full range of other ways that practical reasons can be second-personal.

A further illustration of this general issue can be provided by looking at the differences between two reactive attitudes: blame and resentment. This is my next little case-study. Blame is an attitude that, as Strawson highlighted, implies that the person being blamed was responsible for the action for which they are being blamed. In light of this responsibility, blame is a kind of anger directed specifically at the agent *qua* doer of the deed in question. That is, blame is an intensely focused kind of displeasure to feel towards someone: when A blames B for φ-ing, A is displeased with B precisely insofar as B φ-ed. In being focused in this way, blame *as such* does not extend to a general displeasure or anger towards B themselves. To think of blame as an expression of a judgement that an action was wrong, is to think of blame as a *deontic* reactive-attitude.

Only when blame is construed deontically can it be claimed that blame is twinned with moral wrong-doing in such a way that blame is appropriate if and only if a moral wrong has been committed. And this is indeed claimed, by Darwall (2006, pp. 17-18), who himself cites the support of Strawson, Wallace, Watson, Bennett and Scanlon. So one serious position on the nature of blame has it that blame is an emotional response to a deed, one that is directed towards a person, but only as doer of the deed. If blame as such contained within it a broader negative sentiment towards the agent than just that they had committed some wrong, then it would not correlate in such perfect symmetry with moral wrong-doing. That is, if blame also implied a negative character judgement or some personal dislike of the blamed party, then there would be cases in which an act of moral wrong-doing would not entail the appropriateness of blame. This is because there will be contexts in which someone has done something wrong but that alone does not affect how their moral character is – or ought to be – appraised.

One example to illustrate this might be where Wynnie loves Donna to such a degree that Donna's status in Wynnie's eyes is impervious to the effects of a little wrong-doing. In such a situation it is perfectly intelligible for Wynnie nonetheless to blame Donna for some misdemeanour that harmed a third-party – to believe that Donna is responsible for it, that she shouldn't have done it, and even feel peeved towards Donna for having done it – all without loving Donna any less. If this case is intelligible then it makes sense to blame someone without thereby disvaluing them, so blame cannot necessarily imply such a broader negative valuation of the character of the blamed party. Similarly, suppose that Wynnie despises Donna, thinks of her as the lowest most contemptible scoundrel, that Wynnie expects nothing other than the worst from Donna. It would still be possible, if

Donna were to commit a new act of wrong-doing, for Wynnie to blame Donna for it. Wynnie would be able to add this new entry to the litany of accusations against Donna, and to do this *without* thereby intensifying the contempt which, one can imagine, cannot be intensified any further by the discovery of such a wrong.

Blame, then, is plausibly understood as a purely deontic attitude: an evaluation of a deed done by a person, but not an evaluation of the person themselves. A question I would like to raise is whether the same can be said for resentment. Darwall (2013b, p. 22) – again conscripting Strawson in support – suggests that blame and resentment *both* be understood as deontic attitudes. Where blame is an 'impersonal' reactive attitude that anyone can feel towards anyone else guilty of wrong-doing, resentment is 'personal', felt by one wronged, towards the one who wronged them. On this picture, just like blame, resentment is a pure reaction to a deed, and does not intrinsically involve any broader evaluation of the guilty party.

A problem with this pleasingly neat analysis of this pair of reactive attitudes is that it draws a clean divide between judging someone's actions, and judging them, whereas it is not obvious that such a divide is reflected in our actual attitudes. It is particularly doubtful, I suggest, that resentment is conceptually distinct from the valuing attitudes that flow in either direction in a bipolar relation. If A wrongs B, and B reacts with resentment towards A, the character of B's feeling towards A is, I suggest, inextricable from the dynamics of B's other feelings towards A. It is one thing to suggest that such an attitude as moral blame might be dispassionate in the sense of being detached from the setting of ordinary interpersonal valuing-attitudes, but it is quite a further stretch to say the same about resentment, where the latter attitude is always situated in the relationship between two particular individuals.

The examples of Wynnie and Donna illustrated the intelligibility of blame in the absence of altered interpersonal valuing-attitudes. But the same kind of examples would be jarring and even comparatively *unintelligible* if the wrongs in question were violations of bipolar obligations – that is to say, if Donna's wrongs were done *to* Wynnie. In the first case, suppose that Wynnie loves Donna deeply and Donna wrongs Wynnie in some way. Even if Wynnie swiftly forgives her, and even if her overall attitude towards Donna always remains one of great affection, discovering that Donna has wronged her *automatically* registers in the ledger of Wynnie's appraisal of Donna. To acknowledge a wrong done directly to herself consists partly in reacting negatively towards the wrong-doer. Whereas it might be possible to dispassionately acknowledge a wrong done to another, proper recognition of a wrong done to oneself is in part an emotional response. And interpersonal relationships are such that the currents of interpersonal valuing-attitudes – affection, respect and whatnot – are always affected by the emotional response of resentment.

The same result is obtained in the opposite kind of Wynnie-Donna example, in which Wynnie despises Donna. Here too, it was possible that when Wynnie found a new thing to blame Donna for – a wrong done to some third-party – it did not alter one bit how Wynnie felt towards Donna. Indeed,

she would have expected nothing less. But if the new wrong is inflicted upon Wynnie herself, then recognising it inevitably inflames her dislike for Donna. Each new wrong committed within a bipolar relation will exacerbate the negative sentiments that characterise one's appraisal of the other.

If this is right, and resentment unlike blame is deeply linked to the interpersonal valuing-attitudes which always flow through bipolar relations, then a deontic construal of resentment is inadequate. But it is unsurprising that modern moral philosophers construe such emotions in deontic terms because moral philosophy typically assesses moral attitudes in isolation from the fuzzier realm of interpersonal valuing-attitudes. There are many hues of endearment, pity, esteem, distaste and many other-coloured attitudes through which people see those with whom they interact. So it is understandable to prefer to characterise moral attitudes, as much as possible, separately from any consideration of the basis of the judgements underlying such appraisals. But this separation will always leave out something of the complexity of how we respond to one another's deeds, and thereby to the very conditions of appropriateness of judgments of worthiness for attitudes of a kind with resentment.

Again, this is a matter that would be soluble in a thorough account of the practical significance of the second-person relation. That kind of account would explicate the complex relation between deontic attitudes and evaluations of character as the two are interlaced when we interact in second-person relations with one another. Illuminating whatever commonality there may be between the second-personal character of blame and resentment is a goal of my discussion of Darwall in Chapter Two. A broader issue that this connects to is one that motivates my discussions of love in Chapters Three and Five, namely, the matter of how *personal* (bipolar) relationships figure in determining the appropriateness of moral attitudes like blame which can be expressed in *impersonal* contexts, like between people who are strangers to one another.

A final illustration of the motives of the present project can be seen by inspecting the boundaries between moral obligations pertaining to the lives and ends of others, on the one hand, and on the other, supererogatory reasons pertaining to the same. These boundaries exist wherever there are moral obligations to do with other people – there is always a nearby possible situation (if not an actual situation) where one could go beyond one's obligations to do something morally good for another that could never have fairly been expected of one.

There are a variety of accounts of what characterises an obligation. Joseph Raz (1999, chpt.2) advances an influential account according to which what distinguishes obligations is their *exclusionary* quality: they preclude other reasons that might be present from figuring in the deliberations of an agent. On Raz's view, if one is obliged to fetch a bucket of water, then one's yearning to lie on the grass and watch the clouds is silenced by the obligation – the yearning is excluded from even counting as a reason in one's practical deliberations. Alternatively, with Darwall (2011), one might construe obligations as just those reasons which, if unheeded, license moral blame.

Thus one is obliged to fetch a bucket of water if and only if failure to do so would render it appropriate for others to react with blame.

On either of these two views, obligations are clearly demarcated from supererogatory reasons. They say, respectively, that supererogatory reasons differ from moral obligations in lacking any exclusionary power over other reasons, or in failing to carry the threat of blame in the event of their going unheeded. An interesting feature of this clear-cut theoretical boundary is how blurry it can look in real life. Let's use the term 'altruistic reasons' to refer to all reasons to act in the interests of others (ignoring the self-sacrificing connotations that the term sometimes otherwise has). All altruistic reasons could be plotted on a graph, where one axis represents the extent of the sacrifice required by the subject in order to heed the reason, and the other represents the extent to which acting on the reason will advance the interests of some other person. Some such reasons are obligatory. Others are not. And whether such a reason is obligatory is partly determined by the two factors represented in the graph.
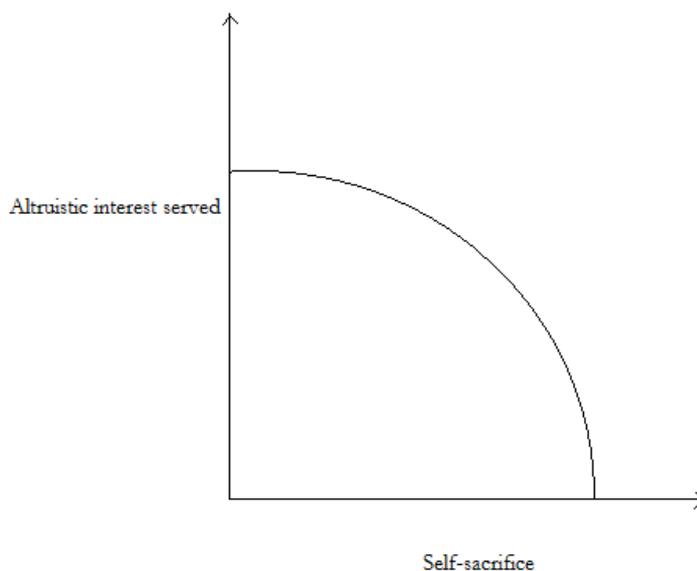
Figure 1

The vertical axis here represents the altruistic interest served but inverted, so that at the base of the axis will be acts which maximally serve the interests of another – like saving someone's life. At the top of this axis will be acts which serve another's interests only a small amount – like, say, washing their dishes. But the lateral axis progresses positively so that at the left end will be acts which constitute a small sacrifice, ranging to the right end where the acts constitute significant sacrifices.

The curve depicted might be thought to mark the boundary between obligatory and supererogatory moral reasons. According to this boundary, any reason that falls within the quarter-circle marked is obligatory because it serves a sufficiently significant altruistic interest at a

sufficiently insignificant personal cost. On the other hand, the altruistic interest served by other actions not encompassed by the curve is still a good thing. Indeed, serving the interests of others is a *morally* good thing. So those potential actions give rise to supererogatory reasons which fall outside of the quarter-circle. These are good things to do, all things being equal, and doing them often warrants moral praise; but they are beyond what could be expected of the agent in question.

Now, the purpose of sketching this boundary is to draw attention to the peculiarity of a theory – such as either Raz's or Darwall's or indeed, any other of the most popular accounts of moral obligations in contemporary moral philosophy – which offers a fundamentally different account of both the grounds and the normative force of the reasons that lie on either side of the curved line. Such theories posit some kind of *authority* as a crucial concept for making sense of why one ought to pay heed to reasons on the inside of the line, though no authority will have any role at all to play in explaining why one ought to follow any of the outer reasons.[6] This does seem peculiar, given that in any real-life predicament that sharp boundary may be very difficult to see – perhaps even impossible.

In fact, the very idea of a firm boundary to one's obligations is infused into certain currents within contemporary culture in such a way that it can be invoked to appease the consciences of people who have the power to help others but who choose instead to help themselves. The notion of a hard boundary itself hardens people against the needs of others – at least it can do. Think here of wealthy individuals who absolve themselves of responsibility for less fortunate people around them with the assurance that having paid their taxes and met the responsibilities of their families and their social roles, they have done what morality expected of them. Such a thought makes perfect moral sense – from an individual perspective – of societies which ought to make no moral sense at all, where unimagined technological and economic power is juxtaposed to large-scale impoverishment and its concomitant levels of suffering.

This is not to say that moral theorists like Darwall and Raz are themselves in any way implicated by the bad-faith pseudo-justifications of powerful people who do what they take their duties to be, and leave it at that. But it is to say that a better conception of the relation between moral obligations and supererogatory reasons would be one which highlighted their shared foundations – their shared basis in the value of the interests of others. An account of the practical significance of the second-person relation could yield such an improved conception of that boundary. Because it is when we relate to others second-personally, as people, when we understand their perspective as one of a subject – much like ourselves – whose ends and desires are intelligible (and indeed evaluable) – it is then that we become acquainted with the normative force of reasons to act in others' interests. Altruistic obligations and supererogatory reasons have a common second-personal character. Better

---

[6] I say 'some kind of authority' because the notions of authority vary widely between different theories of obligation. On Raz's view, for example, one can understand obligations without any supposing any person to possess what one might call *personal* authority. But his picture does nonetheless include *some* kind of authority in the very notion of an exclusionary reason, which has exclusionary authority over other reasons.

understanding this character is a driving motivation of this dissertation, especially my discussion of agapic love in Chapter Five.

3. The plan

With a view to filling these explanatory gaps, the remaining five chapters will take up and address a programme of practical philosophical issues connected to the second-person relation. As noted above, the dominant theory in the moral philosophy discourse on the second-person relation is Darwall's. It makes sense, therefore to begin, in Chapter Two, with an analysis of his outlook. Ambitiously, Darwall offers an account of moral obligation – of the bindingness of moral norms – by appealing to the standing that people have for one another when they occupy the second-person standpoint. Thus, on his view, the very forcefulness of moral reasons is exclusively a practical consequence of second-person relations. In this context, Darwall coins the term 'second-personal reason' to refer to those demands that can be made from the position of a second-person, whose unique pull is derived from the relation between two mutually recognising subjects.

The goal of my analysis in Chapter Two is to give a clear presentation of Darwall's theory of moral obligations *as* second-personal reasons, and then to advance some criticisms of that theory. Specifically, I have two challenges to put forward. The first is that what Darwall refers to as second-personal reasons is in fact only a subset of the larger set of practical reasons whose normative force derive necessarily from a distinctly second-person relation between two agents. The second challenge that I will offer is that those reasons that Darwall *does* identify as second-personal are, in fact, not. That is, I will offer an argument to the effect that Darwall's account of moral obligation fails to ground those obligations in the natural authority of the 'you' (and thereby fails to ground them at all).

Chapter Three builds on the first-half of Chapter Two. Since moral obligations do not alone exhaust the terrain of second-personal reason-giving, I will schematically categorise the families of reasons that do. In this spirit I will consider the different ways in which our relating second-personally to one another might affect, or even ground, reasons of beneficence, of friendship, of love; reasons from explicit contractual understandings, and relatedly, the reasons that arise from favours, expectations, commands, demands and requests. My goal in that chapter is to clarify the senses in which, and extents to which, different practical reasons might be thought to be second-personal. Further, I will investigate and argue against the idea that some reasons might be second-personal 'all the way down'.

Following on from that, in Chapter Four, I will take up a particular focus on the reasons presented in requests. These, I shall contend, are a paradigm of second-personal reason-giving. I will assess some existing attempts to explain the distinctive normative move of requesting, which is an act of intentionally creating a new, discretionary reason for one's addressee. I find, however, that extant accounts do not meet the adequacy conditions for a theory of the normativity of requests, so I offer a

suggested amendment to the most developed account (one put forward by David Enoch). On the amended picture, requests are keyed necessarily to interpersonal valuing-attitudes of the sort that arise especially in second-personal relations.

Such interpersonal valuing-attitudes are a broad family, and characterise the full range of human relationships including, amongst other things, familial affections, romantic love, and agapic love – that is, the love of people as people. In the fifth chapter I hone in on the notion of agape and investigate whether it is itself an attitude held from the third-personal, observer's standpoint, or rather from that of an active, participating, second-person. I argue for a controversial and revisionary conception of agape according to which it is more akin with personal loves (such as romantic love) than is commonly assumed. In this chapter and the next I seek to bring into the discussion some of the historical antecedents to the contemporary philosophy of the second-person. Specifically, in this chapter I will engage with Martin Buber and Simone Weil. Only the former wrote explicitly about the second-person, but both were pioneers of distinctive ways of appreciating the value of other people (as people) that differ from the disengaged standpoint of an observer.

The sixth and final chapter will return to the idea – explored from one angle in Chapter Two – that all moral reasoning might be in some way underpinned by second-person relations. However, rather than renewing a line of thought similar to Darwall's, I take up some insights from Emmanuel Levinas. I explore the possibility that second-personal encounters constitute a rupture in experience that can play a crucial role in moral epistemology: namely, explaining how a subject could become acquainted with an intrinsic value totally outside of itself. I will assess this Levinasian hypothesis from several angles, and offer some observations to support his idea that the face to face encounter with another person can in a sense transcend the conceptual parameters that otherwise mediate experience. If Levinas' view is vindicated, then it follows that moral reasons depend on second-person relations.

Chapter 2
Do second-person relations explain
(only) moral obligations?

Here is a phenomenon worth explaining: other people matter to us. Specifically, the immediate others around us, whether friends or strangers, matter in ways that guide how we should act, and constrain what it is permissible for us to do. Put another way, one could say that for us, as agents, other people have practical significance. This is surely true, but in order to explain *how* others matter, one might need to be more precise. To this end, a distinction can be drawn between two contrasting ways in which one might relate to other people; between two perspectives from which other people appear to guide or constrain one's practical outlook quite differently.[7] I will introduce this distinction again, as a reminder of a key message from the last chapter.

On the one hand, one might relate to others *third-personally*. This mode of relation is that in which one relates to others objectively, rather than as a subject like oneself.[8] That is to say, when viewed from the third-person person perspective, other people may indeed be regarded as valuable, as possessing characteristics in virtue of which people can matter, but they cannot be regarded as partners in any distinctly interpersonal interaction. Paradigmatically, from the third-person perspective others are not partners in dialogue. This means both that one cannot conceive of addressing them with reasons – in the form of requests, commands, demands, advice, offers or anything – and also that one cannot conceive of being addressed by these others with reasons oneself.

On the other hand, of course, one can relate to others *second-personally*. This is just that perspective from which the other person is regarded as a partner in distinctly interpersonal interaction. It is the way we think of someone when we address them, when it is appropriate use the pronoun 'you'. As such, engaging with another second-personally means entertaining the possibilities both of addressing the other with reasons, and of being addressed by the other with such reasons oneself.[9]

Drawing this distinction does not imply that the second- and third-personal perspectives cannot overlap, that one could not relate to another at once as both an object (with certain objective traits), and as a subject and a partner in dialogue. One can. But the conceptual distinction enables the following question to be raised: are there reasons for action that arise only from the second-person perspective?

---

[7] The idea of this distinction has caused recent excitement amongst philosophers with some heralding a 'paradigm shift in the way we address questions about a range of fundamental issues in these fields' (Eilan 2014, p. 265).

[8] This formulation of the nature of the third-person relation is not perfect because there is a sense in which one can acknowledge a person's subjectivity from an objective, or side-on, or observer standpoint. So in order to understand what third-personality is *not*, one must construe 'relating to a person as a subject' as more than merely acknowledging that a person *is* a subject. Relating to a person as a subject is *engaging with them* as a subject, and as such it is neither a side-on, nor an objective, standpoint; as such, it cannot be third-personal.

[9] See (Haddock 2014) and (Rödl 2014).

As mentioned in the previous chapter, this question has become prominent in contemporary practical philosophy thanks to Stephen Darwall and especially his book *The Second-Person Standpoint* (Darwall 2006, *SPS* hereafter). He answers affirmatively with an account of 'second-personal reasons' as the conceptual basis for all of morality. In this chapter, I will argue that Darwall's theory does not explain the practical significance of the second-person. Firstly I will illustrate that the theory is too narrow to provide a full explanation of that practical significance.[10] It is too narrow because it overlooks the normative significance of any given *particular* other, focusing, as it does, on the normative role played by others *in general*. And secondly, even within its narrow domain, I will press a concern over whether Darwall's view succeeds in its own terms to adequately explain even this general practical significance of the second-person.

The structure of the chapter will be as follows Section one will give a precis of Darwall's theory of morality as a system of second-personal reasons, emphasising the way that Darwall can seem to purport to account for the practical force that a person can legitimately exert over another in the course of real interpersonal interaction. In section two I will raise the first of my criticisms, and show that in fact the normative grounds of what Darwall thinks of as a second-personal reason are completely independent of any consideration about any particular other person. The point here is that Darwall's account of the practical significance of the second-person standpoint is incomplete in that it omits any explanation of the practical significance borne by particular other second-persons. This conclusion that I will reach in section two opens the way for later chapters of this dissertation, which will examine the set of phenomena that Darwall overlooks: the practical significance that particular others bear in second-personal interaction. Putting aside his omissions, section three will return to the aspect of second-personal normativity that Darwall's theory does encompass. There I will level an objection against the success of his account.

1. Second-personal address as the foundation of morality

> Morality involves a distinctive kind of accountability by its *very nature*. If I fail to act as I am morally required without adequate excuse, then distinctively second-personal responses like blame and guilt are thereby warranted. (*SPS*, pp. 26-7, emphasis added)

---

[10] An important caveat to make here is that in his more recent work since *SPS*, Darwall has contributed significantly to discussions of the broader significance of the second-person. Moving beyond the scope of second-personal moral obligations that are the focus of *SPS*, he has written on the role of the second-person relation in, for instance, relations of trust (Darwall 2017) and of love (Darwall 2016). So, the first of my two critical points in this chapter is not that Darwall himself is confused about the true breadth of the practical significance of the second-person: evidently he agrees that second-persons affect our practical reasoning in more ways than just shaping the moral obligations that constrain us. Rather, the target of my argument – as will become clear – is an implicit current in *SPS* that suggests that the felt force that particular second-persons can exert on our practical reasoning is something that *SPS* accounts for. I will argue, on the contrary, that such felt force of particular others is an omission – it falls outside the scope of Darwall's purview in that book.

This is the central idea that Darwall advances in *SPS*: that not only do we owe it to each other to act morally, but that it is part of the very nature of morality that people hold one another accountable for their moral conduct. This basic conviction is manifested in four key concepts that Darwall thinks form an 'interdefinable circle' and that in concert can be used to explain the nature of morality (*SPS*, p. 12). They are: accountability, second-personal authority, a second-personal reason, and second-personal address.[11] As a means of presenting Darwall's account, I will explain these core concepts and set out how each is thought to entail the others. What I hope to make clear in this portrayal is that part of the appeal of the moral outlook that Darwall articulates, is that it captures a striking feature of our moral lives, one that is not otherwise given such prominence in moral theorising: namely, the significance that people have for one another as both sources and arbiters of practical reasons.

I have already raised the question of whether there are reasons for action that arise only from the second-person perspective. Darwall's answer is that there are, and they are moral obligations.[12] As suggested by the quoted passage above, the grounds for construing morality in terms of second-personal relations is the notion of accountability and its conceptual connection with the notion of responsibility. Darwall claims there to be an analytic conceptual connection between responsibility and accountability such that to be responsible for doing something just is being responsible *to* someone for doing it, and being responsible to someone just is being in principle accountable to them (*SPS*, p. 67).

To bolster this connection, Darwall enlists an insight from Mill, to the effect that moral wrongness already implies blameworthiness (*SPS*, p. 27; Mill 1998, chpt. 5). To judge something to be wrong always necessarily means that it warrants others in blaming the wrong-doer, and warrants the wrong-doer in feeling guilt (unless they have a valid excuse).[13] This insight is important because in coming to recognise that moral transgressions are always the business of the moral community, we come also to see that members of the moral community have a certain kind of standing for one another, the standing – or authority – to hold one another to moral account: 'When we are morally obligated, we are not free to act otherwise; members of the moral community have the authority to hold us responsible if we do' (*SPS*, p. 27).[14] It is on this basis that Darwall's circle of core concepts includes responsibility which entails and is entailed by accountability, which entails and is entailed by second-personal authority.

---

[11] See *SPS*, p. 12.

[12] Strictly speaking, Darwall thinks that whilst all moral reasons are second-personal, not all second-personal reasons are moral. For instance, he thinks of both a sergeant commanding her troops (*SPS*, p. 60), and a diner ordering some food of a waiter (*SPS*, p. 51) as instances of second-personal reasons.

[13] Or, as Darwall puts it quite directly elsewhere: 'What it is [...] for an action to be morally obligatory and its omission morally wrong, is just for it to be an action the omission of which would warrant blame and feelings of guilt, were the agent to omit the action without excuse' (Darwall 2013a, p. 177).

[14] It is worth pointing out that Darwall's view is not that absolutely anyone always has the standing to explicitly confront anyone else and reproach them for a given moral transgression, but rather just that absolutely anyone can at least legitimately blame anyone else for such a transgression, where such blame *implies* an authoritative demand (Darwall 2013b, p. xiv n.1).

To clarify the last step there: the inclusion of the concept of authority is again an a priori claim. Namely, from the fact that when someone is guilty of a wrong anyone is entitled to blame them, it is deduced that people generally – just by virtue of being members of the moral community – have the authority to hold one another to account. A premise of this deduction is that blaming someone *is* holding them to account. The idea of blaming as a mechanism of accountability introduces a further concept of the circle, that of second-personal address.

This is the component of Darwall's theory that brings this picture of moral responsibility to bear on the question that I raised above, of whether there are any reasons that arise only in the context of second-person relations. Darwall thinks that moral reasons are demands that could be made of agents by others who occupy the standpoint of the second-person and express those demands through a distinctly second-personal form of address. What makes it distinct, Darwall thinks, is that when we address another as 'you' and present them with a reason for action, telling them what we think they should do, we presuppose that they are free and rational beings, and we communicate that presupposition to them through the address.[15]

This conception of second-personal reason-giving has some compelling features. Firstly, Darwall notes that when demanding of someone that they meet their moral duty, in order to take oneself to be *rationally* guiding their will, one must suppose that the addressee has the rational capacity to recognise the normative fact that they have a duty (and that it is binding). Moreover, they must be the kind of free agent who has the further capacity to act on reasons that they take to be normatively binding. Darwall goes on to endorse a further presupposition – one he draws from Fichte. Namely: that to address the other with a demand to meet their moral duty supposes that, being free and rational as they are, the addressee is the sort of person who *bears* moral duties; that is, that they are a responsible member of the moral community. Given that the addressor is assuming her own standing to make a demand of the other, she must also assume herself to be, similarly, an equally responsible member of that same moral community.

If all of these presuppositions are satisfied, then when one addresses another with a demand that they fulfil their moral duty – either explicitly, or implicitly through a reactive attitude like blame – one gives a second-personal reason. The idea of a second-personal reason, then, is the final piece of the circle of irreducible second-personal concepts.[16] With this system of ideas, Darwall provides a way of understanding the practical significance that other people have that is distinctive of second-personal engagement with those others. Namely, this: when one relates to others second-personally – and only then – one recognises their status as equally deserving of the respect due to responsible

---

[15] 'To enter intelligibly into the second-person stance and make claims and demands of one another at all, I argue, you and I must presuppose that we share a common second-personal authority, competence, and responsibility simply as free and rational agents' (*SPS*, p. 5).

[16] 'A second-personal reason is one whose validity depends on presupposed authority and accountability relations between persons and, therefore, on the possibility of the reason's being addressed person-to-person' (*SPS*, p. 8).

members the moral community. As such, in engaging with others second-personally one is confronted with a set of reasons that constrain what one could legitimately do, by one's own lights. These are second-personal reasons, moral responsibilities, that other people (second-persons) have the right to demand, and to which others have the authority to hold one accountable.

So goes Darwall's answer to the question of what practical significance other people have when we relate to them second-personally. However, I think it is incomplete. That is, I will now argue that Darwall's view only goes as far as explaining what practical significance a given second-person might have in terms of the *general* possibility of there being second-personal relations at all. But one might think that a given particular second-person might possibly have further practical significance – a practical force – that might be particular to the address that they make of a given subject (or first-person). Indeed, much of the rest of this dissertation will explore that practical significance that others might wield within particular second-person relations.

I also think that Darwall's answer is, even as far as it goes, inadequate as an explanation of even the general role played by second-persons in practical reasoning, but I will come onto that in section three.

2. An oversight: the missing role of individual authority

Darwall's flagship example of a second-personal reason is one addressed by someone whose foot is being trodden on, to the person treading on it. Of that reason, Darwall says this:

> What is important for our purposes is that someone can sensibly accept this second-personal reason for moving his foot, one embodied in your claim or demand, only if he also accepts *your* authority to demand this of him (*SPS*, p. 8, emphasis added).[17]

If this claim were true, then Darwall's picture would be one that recognises – and indeed gives central prominence to – the practical significance that a particular other has when one relates to that other second-personally. The account would succeed in recognising that significance because it would be conceiving of the normative grounds of second-personal reasons as stemming directly from the authority of the particular person who addresses the reason, the addresser. And having a special authority that uniquely grounds reasons for others to act a certain way would indeed be having a special kind of second-personal practical significance.

But in fact the claim in this quoted passage is misleading.[18] Darwall is not entitled to attribute this normative significance to the authority of the particular individual with whom one is engaged in

---

[17] This is not a one-off slip of the tongue. Darwall implies that the authority of a particular other is normatively significant at various other points in the book. Here is another example: 'If, however, you see his taking his foot off yours to be dictated by mutual respect, you will see him as having a reason that is grounded in *your* authority to demand that he move his foot (indeed, to demand that he not have stepped on yours in the first place)' (*SPS*, p. 129, emphasis added).

[18] Again – as in footnote 10 above – I must qualify this allegation. I do not mean that Darwall himself deliberately misleads his readers into thinking that his claims are more relevant to considerations of particular others than those claims actually are. Rather, my point is that one *could* read *SPS* and be misled in this way into

interaction. In his theory, it is not true that one can sensibly accept second-personal reasons only if one grants authority to the particular person who addresses one with a demand. In his theory, one could sensibly accept such a reason without recognising the authority of the particular addressor, so long as one could reflectively imagine a legitimate authority making the same demand in the same situation. As such, there is no role in Darwall's picture for the authority of the actual individuals who address one another with demands, and elicit from one another the range of familiarly second-personal emotions. To demonstrate this, I will examine how Darwall might ground this claim about the significance of a particular other, and show that these grounds are incompatible with deeper commitments that underscore his outlook.

Perhaps the most natural way of understanding Darwall's apparent claim – that one can only sensibly accept a second-personal reason if one recognises the authority of a particular addressee – is as implying the following bi-conditional:

Individual Authority View: A has a second-personal reason to φ if and only if B, qua particular individual, has the authority to demand of A that she φ.

The Individual Authority View might well be thought to be implied by claims such as the one quoted at the beginning of this section. What's more – to reiterate – if this view is true, then it constitutes a substantive affirmative answer to this chapter's driving question of whether there are reasons for action that arise only from the second-person perspective. Namely, the reason that A has to φ – A's moral responsibility – arises only in virtue of A's relating second-personally to a particular individual, that is, B.

The problem with the Individual Authority View from Darwall's perspective is, in short, that it is incompatible with the idea of second-personal reasons as things that all moral agents owe to *all other* moral agents (and not just to a given particular agent). This is a fact that Darwall is quite conscious of. In fact, in a paper responding to challenges raised by R. Jay Wallace (2007), he explicitly rejects the Individual Authority View as an account of second-personal reasons (Darwall, 2013b, pp. 35-39). In that paper, Darwall distinguishes a sub-set of second-personal reasons that he calls *bipolar obligations*,[19] which falls within the broader category of second-personal reasons that all moral agents can demand of one another, which he terms *moral obligation period*. The distinction between the role of the particular individual's authority in grounding the two kinds of reasons is telling.

A bipolar obligation is a reason for a particular individual (the obligee), that stems from the authority of another particular individual (the obligor).[20] Thus, with regard to bipolar obligations, the

---

thinking that Darwall ties the bindingness of morality to the pressure that people wield over one another in particular interactions between individuals.

[19] Following the important discussion of bipolar normativity in (Thompson 2004).

[20] Somewhat confusingly Darwall uses these same terms, obligor and obligee, but the other way around so that the one with the obligation is the obligor (Darwall 2013b, p. 20). I think the motivation there is for the terminology to reflect the distinction between a promisor and a promise, where again it is the promis*or* who has

Individual Authority View is true. That is, in these cases the agent in question *only* has a bipolar obligation if a certain particular individual – namely an obligor – has the authority to obligate that agent. The paradigm cases that concern theorists of bipolar normativity – such as Wallace and Thompson – are things like promises, where the reason you have to keep your promise clearly has something very much to do with the wrong that reneging on the promise would do to the promisee. In such cases, the particular other – the obligor – has a special standing above and beyond that of any other member of the moral community. They have this authority qua the particular individual to whom the promise was made. This is true in such cases despite it also being true that since we have a moral obligation in general to keep our promises, anyone at all (by virtue of being someone) also has some standing to blame the offender for reneging on their promise.

However, not all second-personal reasons are like promises. What goes for bipolar reasons in particular, does not go for moral obligation period, which is to say, it does not go for what Darwall thinks of as second-personal reasons in general (Darwall, 2013b, p. 36). For the broader category of second-personal reasons, no particular individual necessarily occupies the status of a unique obligor. As such, the Individual Authority View is false as an account of second-personal reasons in general. In the case of second-personal reasons in general, no particular individual's authority could possibly be necessary for such a reason to exist, because such reasons stem from the authority that *anyone* would have to blame the agent for non-compliance. In the terms Darwall uses in the same essay, second-personal reasons are grounded in *representative* authority, rather than individual authority (Darwall, 2013b, p. 23).

It is important to recognise that the insistence that second-personal reasons are grounded in representative authority is not an *ad hoc* way of distinguishing such reasons from the challenge posed by theorists who stress the significance of bipolar obligations. On the contrary, the idea of representative authority is at the very heart of Darwall's idea of the second-person standpoint. In the schema of *SPS*, the importance of representative authority is manifested in the part that Darwall calls 'Pufendorf's Point':

> genuine obligations can result only from an address that presupposes an addressee's second-personal competence. To intelligibly hold someone responsible, we must assume that she can hold herself responsible in her own reasoning and thought. And to do that, she must be able to take up a second-person standpoint on herself and make and acknowledge demands of herself from that point of view. (*SPS*, pp. 23)

The very idea of addressing someone second-personally, on Darwall's picture, just is providing them with reasons the bindingness of which they themselves assent to by entering into the social world of second-personal interaction. That is: when one makes a second-personal demand, one is not attempting to pressure the other into some deed *just* because one expects it of them. A demand does indeed express that the addresser expects the action of the addressee, but second-personal demands

---

the obligation. But in this dissertation I will stick with what I think to be a more natural use of the words: the obligee, like an addressee (a requestee, a commandee, etc) is an agent facing a reason, and an obligor, like an addressor (requester, commander), is one giving a reason.

also express that the addressee ought to expect this action *of themselves*. This is the crucial point. Any given agent can expect it of themselves – and demand it of themselves – that they comply with all second-personal reasons (that is, of course, all those that are agent-relative to them). Such self-expectation is nothing but an inward reflection of the expectation of the moral community at large. Thus, since second-personal reasons always presume that addressees could hold themselves to account for compliance, such reasons must be grounded in the representative authority of the moral community.

Given, then, that second-personal reasons must be grounded in this more universal species of authority, it will be worthwhile to clarify exactly how this is incompatible with the Individual Authority View as I labelled it earlier. To do this, I shall define the competing conception of authority and illustrate the contrast between the two views through two examples. One example will demonstrate that there are such things as second-personal reasons for which there is no unique obligor whose authority to demand it is the ground of the reason. The second example will show that even in cases of reasons that are both bipolar obligations – with a uniquely authoritative obligor – and moral obligations period, there is a clear conceptual distinction between the two. That conceptual distinction reveals that insofar as a reason is merely second-personal in Darwall's sense, it is not grounded in the authority of any particular individual.

In light of the preceding discussion, here is a characterisation of second-personal reasons that Darwall is quite fundamentally committed to:

Representative Authority View: A has a second-personal reason to φ if and only if any member of the moral community has the authority to demand of A that she φ.

My question is whether the Representative Authority View can account for the practical significance of second-person relations. A proponent of the view might argue that they can provide such an account, by claiming that the Representative Authority View is compatible with the Individual Authority View. But the views are in fact incompatible.

To see this, consider first the case of Bonnie who is a somewhat reckless inventor and self-made scientist. Bonnie has her mind set on undertaking a set of experiments that are likely to cause quite large explosions. She couldn't possibly carry out the tests enclosed in her laboratory at home, so she decides to take her equipment outside – and drives to what she thinks to be a suitable spot in a clearing in a forest, many miles from the nearest human dwelling. Now, suppose that one particular experiment is sure to cause a very large explosion and it will quite possibly start a fire in the woods. Bonnie realises this possibility and she knows that the forest is home to deer, numerous smaller mammals, birds, amphibians and uncountable insects.

I think it is clear that Bonnie has a moral responsibility not to conduct this dangerous experiment. As such, Darwall's theory of morality has it that Bonnie has a second-personal reason to abort the experiment. That is, a reason that exists because, and only because, any member of the moral community – including herself – has the right to demand that she abort. Any person also has the

authority to hold her to that demand, and to blame her if she goes ahead with it. To be sure, there are other, third-personal, reasons to abort the experiment; but Darwall's Representative Authority View says that there is also a reason that stems from the possibility of members of the moral community holding one another to account for their shared moral commitments.

Since, *ex hypothesi*, there is no individual person who is a potential victim of Bonnie's shenanigans there is therefore no uniquely situated obligor with a special claim right to personally demand that Bonnie aborts the experiment, or with the standing to *personally* resent her for her action. Thus, Bonnie is not under a bipolar obligation to anybody. To put it in terms of the Individual Authority View, there is not necessarily any particular individual, any one specific person B, whose authority over A generates A's responsibility. The Individual Authority View therefore has it that there is no second-personal reason. Clearly, then, the Individual and the Representative Authority Views are incompatible. But I hope that beyond merely demonstrating this incompatibility, this example elucidates the contrast in the normative grounds that are thought to lie beneath second-personal reasons on the two Views. Whereas the Individual Authority View generates the normativity of second-personal reasons from the practical significance of a particular person, the Representative Authority View makes no essential reference to particular people.

This contrast should cast some light back on how Darwall's theory explains more straightforwardly interpersonal cases. Let's return to the unfortunate case of the person whose foot is being trodden on. Call her Masha. And call the offender Natasha. Here, Masha has a claim right against Natasha. She has the unique standing of an obligor to generate a bipolar obligation to Natasha, and justly to feel resentment. Later, she will also be uniquely situated with the authority to forgive Natasha if she so wishes. The bipolar obligation is generated out of the way that from Natasha's perspective, Masha should matter. The Individual Authority View would have it that it is the particular practical significance borne by Masha (and by Masha alone) that grounds the second-personal reason that currently grips Natasha.

It would be an attraction of Darwall's account of second-personal reasons if it were also able to capture the particular significance that Masha alone bears to Natasha in this instance. Indeed, there are points in the exposition of his account at which Darwall appears to suggest that his view does attribute special status to Masha. Recall the claim quoted above, that the addressee faces a second-personal reason in the foot-treading case 'only if [s]he also accepts *your* authority to demand this' of her. But this suggested promise cannot be redeemed by Darwall's account. On the Representative Authority View, to which Darwall is deeply committed, the kind of second-personal reason that faces Natasha is identical in form to that which faced Bonnie. The reason is not grounded in the authority of any particular individual but rather in that of any *possible* second-person.

In this section I have offered an overview of the dialectic between Darwall's conception of a second-personal reason and the idea of a bipolar obligation as championed by the likes of Wallace and Thompson. In light of that contrast, I have shown that despite initial appearances, Darwall's theory of

second-personal reasons is silent on the practical significance that particular individuals bear in second-personal interaction. In later chapters of this dissertation I will take that silence as a calling: an invitation to expand the analysis of the second-person standpoint so that it encompasses the normativity of relations between particular individuals. The case of Masha and Natasha already suggests the intuitive appeal of the idea that particular others can wield a certain kind of practical significance only within distinctly second-personal interaction. But the practical importance of the second-person relation also goes beyond the sphere of moral demands. I will go on to consider the ways in which distinctly second-personal relations give rise to further practical reasons, both by facilitating certain interpersonal valuing-attitudes such as affection, and by facilitating certain kinds of non-moral reason-giving such as requesting.

3. Imagined second-persons and imaginary normativity?

In what has been said so far, I have claimed that Darwall sometimes seems to present his theory as accounting for the domain of the practical significance of second-persons, and that there is a significant set of phenomena within that domain that *SPS* in fact does not account for at all. But I have also so far been assuming that his account is successful as far as it goes, as an explanation of the normativity of moral obligations in terms of second-personal accountability. At this point, I would like to call that assumption into question. Developing points that emerged in the previous sections above, I will raise an objection to the efficacy of the Representative Authority View in its own terms. To anticipate, the objection is that in the absence of any appeal to the immediate, inherent value of concrete, particular second-persons, Darwall's (so-called) second-personal reasons collapse into a form of either ordinary epistemic reason-giving, or coercion. As I will suggest, this objection can be avoided by a rival account, which, again, will be elaborated in later chapters of the present dissertation.

As has been stressed already, the kind of second-person whose authority, on Darwall's view, generates second-personal reasons is not necessarily a particular second-person. Rather, it is simply the *possibility* of a second-person. As Darwall puts it:

> A second-personal reason is one whose validity depends on presupposed authority and accountability relations between persons, and, therefore, on the *possibility* of the reason's being addressed person to person (*SPS*, p. 8, emphasis added).

A natural question in response is this: what makes this possibility significant? It has been established above that whatever is significant about this possibility, it cannot depend on the practical significance of particular individuals as sources of normative authority. Darwall's answer is that second-personal reasons are 'practically directed' in that they make 'a claim or demand on the will' of an agent (*SPS*, p. 10). But understanding what this answer means requires distinguishing it from three things that it *cannot* mean, if it is going to be satisfactory. In order, I will spell out the dividing lines that mark second-personal address from coercion, advice and commands. This is a challenging task because none of these dividing lines are quite where one might initially expect to find them.

Firstly, the practical force of second-personal address must be kept distinct from coercive force. This might seem obvious, given what was said of Darwall's view in the first section above about the freedom and rationality that are presuppositions of second-personal address. Since the address is one that purports to present the addressee with a normative reason, appealing to them to recognise its bindingness and to act on it of their own volition, it follows that the address is very far away from being coercive. However, matters are complicated slightly by what Darwall says about the sanctions that can be concomitant with the authority to issue second-personal address. On Darwall's view, it is possible that the authority to make a demand could entail (in context-sensitive ways) the equally legitimate authority to 'sanction' the addressee for failure to comply with the demand. His example is of the authority that the person whose foot is being trodden on would have, not only to express blame, but to remove the offending foot by force (*SPS*, p. 51). Darwall does not specify exactly what kinds of sanctions might be licensed by the authority to hold people to moral account. But what the book does say is that if you have the authority to threaten such a sanction, 'then however directive your demand, this threatening notice will not amount to coercion' (*SPS*, pp. 51-2).

What is important, though, is that whatever threats of sanction might be implied by an utterance of second-personal address – physical, emotional, institutional, or otherwise – the consideration of these threats *as such* is not what ought to motivate the addressee to comply with the address. That is, when a second-personal demand is addressed, it is not the threat of legitimate sanction that alters the balance of reason for the addressee in their deliberations. What ought to motivate the addressee to adhere to a second-personal reason in any given case is the deontic fact that they owe compliance to the addresser (by virtue of the fact that the addresser represents the moral community at large). Darwall's project represents an ambitious attempt to explain the legitimacy of punitive authority and the bindingness of moral reasons in one blow. But regardless of whether he is successful in grounding punitive authority, the threat of punishment cannot itself constitute the practicality of second-personal address. Despite the apparent complication, second-personal address is reason-giving in a way that must be entirely separate from the threat of coercion.[21] So when seeking to understand how the mere *possibility* of second-personal address could be thought to generate the normativity of moral obligations, one must look beyond the candidate answer that it is the possibility of legitimate forceful coercion that putatively grounds moral reasons.

If second-personal address is not to be understood as an active threat, but rather, strictly, as the presentation of and appeal to a normative reason, then a natural way to understand the act of address would be as a kind of *telling* or epistemic reason-giving. That is, one might well understand the legitimate, non-coercive dimension of the address to be an attempt to reveal a fact to the addressee:

---

[21] Not all coercive reason-giving is punitive. If at the same time as making her moral demand Masha was seducing Natasha, or bribing her, or hypnotising her, then these might all equally constitute coercive reason-giving. Perhaps there are contexts in which these forms of non-rational reason-giving are legitimate in the same way that Darwall thinks that some threats are legitimate. But in any case, none of them play a role in explaining the practicality of pure second-personal reason-giving.

namely, the normative fact that they have a responsibility to the moral community to perform a certain action. If this were so then second-personal address would be tantamount to a form of advice. Darwall is explicit in insisting on a difference between full-blown second-personal reason-giving and advice:

> It presumes to tell another person, not just what to do in the way advice does, but also, in some way or to some extent, to do it (*SPS*, p. 49).

This answer, though, is somewhat unsatisfactory. The question we are pursuing is exactly what this 'way' *is* by which the address can be thought to legitimately go beyond merely telling the addressee what they should do. To be clear, despite Darwall's insistence on the distinction, this is a risk for his theory. That is, it is a salient risk for the theory that if no more substantial conception of second-personal address is forthcoming, then it looks as though it *is* best understood as an epistemic kind of address. As has just been shown, in order to mark legitimate address apart from illegitimate coercion, the theory appealed to the fact that it is a normative fact that is being *freely* acted upon by a subject who binds *herself* to the second-personal reason that is addressed to her. For this reason, marking the difference between legitimate address and advice is a challenge. But it is a challenge that is fundamental to the project of *SPS*.

If the address were to be merely epistemic, then it seems implausible that the very *possibility* of this kind of address could generate the bindingness of moral reasons. If all there was to second-personal reason-giving was the passing of information between people about the duties they owe to one another, then that would be one thing. But since Darwall's view is also, as discussed, that those duties are themselves grounded in the possibility of such (epistemic) address, this candidate answer seems totally inadequate. As an explanation of moral address, it does not conceivably underwrite the normativity of moral reasons in the way that Darwall's system requires of it. There is nothing about the possibility of giving advice that appears to generate moral obligations.

I have argued that there is a sense in which Darwall's conception of second-personal address seems as though it involves the threat of coercion, and an alternative sense in which it seems to be reducible to an epistemic kind of reason-giving. But it is clear that neither of these will do the job that is required of second-personal address by Darwall's account. The very possibility of such address must provide the normative ground of second-personal reasons, and therefore of moral reasons. With this in mind, the final concept from which such address must be distinguished is the concept of a command.

A second-personal reason is a demand. And demands are distinct from commands. Where commands deploy the authority of an addresser over their addressee to create a *new* reason, demands are statements of reasons that exist independently of their being demanded.[22] An unfortunate though interesting feature of *SPS* is that the distinction between demands and commands is sometimes blurred. (Not unlike the distinction between demands and threats, and between demands and advice.)

---

[22] See (Postema 2001) for discussions of the role of authority in commands and (Enoch 2014) for a helpful exploration of these different forms of reason-giving.

Darwall thinks that commands are instances of second-personal reason-giving 'in a sense' (*SPS*, p. 8). This is part of what blurs the matter, because elsewhere the term 'second-personal reason' is used to refer to all of and only the set of moral responsibilities that people owe to one another. This confusion has led some commentators to wonder whether Darwall thinks that moral responsibilities are *created* by being commanded – that is, by being uttered by particular individuals through second-personal address.[23] And one can appreciate the temptation of this line of interpretation. After all, Darwall devotes a substantial portion of *SPS* to a discussion of divine command theories (*SPS*, esp. pp. 107-111). His theory is not a divine command view, but it is in drawing directly from that school that Darwall inherits his conception of authority and accountability. For Darwall, second-personal authority is directly analogous to God's authority. Just as on a divine command view one must act morally because God has the authority to command it, so, on the Representative Authority View, one must act morally because any member of the moral community would have the authority to *demand* it. Given this close relationship between the two accounts, it is hard to see exactly how much of the normative structure that is specific to creating a reason through a command has bled through into Darwall's picture of moral demands.

But what is much more clear is that whatever else one thinks about moral responsibilities, they must be in force independently of the contingent event of their being spoken by human lips. Even if Darwall's Representative Authority View is not an accurate description of the moral reason that Bonnie has not to burn down the woods, she does surely have a moral responsibility of some sort nonetheless. And this is irrespective of whether anyone tells her to abandon her wild experiments. So second-personal address, since it addresses moral reasons, cannot be a kind of command. It cannot be the act of speaking a moral reason that creates that reason.

Where does this leave Darwall's claim that second-personal reasons are grounded in the possibility of second-personal address? Darwall elaborates that what makes the possibility of second-address so significant is its practical directedness, the way that such address amounts to actively *telling* the addressee *to* act. I have presented arguments to the effect that Darwall's account of moral normativity fails if second-personal address were to collapse into a form of coercion, advising, or commanding. And yet, besides these possibilities I do not think there is an answer put forward by Darwall. Not only that, but it is not clear that there is any conceptual space at all between the Scylla of coercing and the Charybdis of advising. In pressing this dilemma, I do not take myself to be making an especially novel point.[24] Indeed, Darwall himself is quite aware of exactly this worry (Darwall, 2013b, pp. 37-8). But I hope that the preceding discussion has made clear just how unclear Darwall's

---

[23] See (Wallace 2007, p. 26). Douglas Lavin (2008) also notes this same possible confusion in reading Darwall: 'the link between the moral 'must' and accountability to another is not limited to obligation created by the performance of a particular act on a particular occasion -- some demands are simply "in force." (There is a very strong temptation to read Darwall otherwise -- as holding that all second-personal reasons must be grounded somehow in actual mental acts of making and acknowledging claims and demands -- but I think it would be contrary to his intentions to do so.)'

[24] A similar point is pressed against Darwall by Bob Stern (2019) to whom I am indebted.

own answer is to this fundamental question. The worst case scenario for Darwall's position that could fall out of my remarks here, would be that all there is to the supposedly distinctive practicality of second-personal address is an obscure triangulation between three discrete answers, none of which will do the job.

Conclusion

This chapter has explored the idea that the second-person standpoint might be a perspective from which others have a certain practical significance – a status as the source of reasons – that they do not have from the third-person standpoint. I gave an overview of Stephen Darwall's theory of moral reasons as conceptually dependent on the possibility of second-personal address. However, I found that this theory does not even purport to account for a substantial array of phenomena that should be thought to come under the practical significance of second-persons. In particular, I found that Darwall's theory does not account for the reasons that stem exclusively from *particular* second-persons, rather than the imagined possibility of second-persons in general. Subsequently, I argued that even in his own terms, Darwall's moral picture faces a fundamental challenge. In order for that moral picture to hold up, it must distinguish the special practical force of legitimate second-personal address from a triad of alternative forms of reason-giving: commanding, coercion, and advising. But it is not clear that the vision of morality in *SPS* has the resources to do this.

Chapter 3

Deep second-personality and its limits

> The single man *in isolation* possesses in himself the *essence* of man neither as a *moral* nor as a *thinking* being. *The essence* of man is contained only in the community, in the *unity of man with man – a* unity, however, that rests on the *reality* of the *distinction* between 'I' and 'You'. (Feuerbach 1972, §59)

In the previous chapter I argued that Stephen Darwall's theory of moral obligation in terms of second-personal accountability does not explain the full scope of the impact that second-person relations have on practical reasoning. As Feuerbach proclaimed in the passage above (written in 1843), there is a pervasive extent to which the essence of human practical reasoning is shaped by social relations – the unity of man with man – where the second-person relation can be seen to play a fundamental role in giving rise to that mutually determining social unity. It is with a view to expanding on the germ of Feuerbach's thought here that I will now turn to focus on that remaining range of practical implications, which are left out by focusing merely on deontic second-personality. In the first section of this chapter, I will sketch a taxonomy detailing the variety of further ways that second-person relations might be thought to figure in determining how people should act. In elaborating on this taxonomy, I will discuss the plausibility in each case of the view that distinctly second-personal relations play a necessary role in grounding legitimate practical reasons. Articulating such a schema will clarify several different senses in which practical reasons can be second-personal. It will also become clear that among these different kinds of second-personal reasons, there are some that are – at least putatively – *more* second-personal than others.

Exploring the idea of such a graded scale, in the later sections, I will spell out the notion of *deep second-personality*. Roughly, this is the view that some non-deontic reasons could be second-personal 'all the way down', which is to say that such reasons are irreducibly second-personal: they cannot be completely explained as reasons without reference to some distinctly second-personal relation between agents. Importantly, the kind of irreducible second-personality that I will discuss here is quite different from another kind of irreducible second-personality. As was discussed in the previous chapter, Stephen Darwall also thinks of some practical reasons as being second-personal all the way down, in this sense. Darwall follows Fichte in characterising the second-person perspective as an in-built feature of morally mature practical reasoning. The second-person perspective is, on his view, part of the very structure of agential decision-making. In contrast to Darwall's Fichtean approach, the views that I will discuss in this chapter do not posit the 'second-person within' as a structural feature of practical reasoning. Rather, the views at issue here will offer more local and specific accounts of certain domains of reasons, where they find reference to some second-person relation to be an ineliminable feature of the normativity of those reasons. I will consider two differing

(though in fact compatible) attempts that have been made to vindicate deep second-personality in this sense, and I will argue that both attempts face significant challenges.

The first such attempt pertains to reasons for love.[25] Some theorists think of love as an attitude that is not only rational, but one whose justification must include reference to properties that the beloved bears only in virtue of their relating second-personally to the loving agent. For example, this view might contend that Seren should love Mair because of Mair's love-worthiness-for-Seren, where Mair's love-worthiness-for-Seren cannot be reduced to two distinct components: a normative and impersonal component of Mair's love-worthiness *simpliciter*, accompanied by some descriptive component specifying Seren's fittingness for Mair's love, or Seren's receptiveness to Mair's love-worthiness. Rather, the deeply second-personal view under inspection denies the possibility of such a reduction. It construes love-worthy properties as irreducibly keyed to particular others in the contexts of specific second-personal relations with those others. If this is right, then it will follow that not only the reasons one thus has to love one's beloved, but also a host of reasons to undertake further actions, which arise from one's loving relationships, are deeply second-personal. For ease of reference, I will label this view *deep second-personality of the other*.

The second attempt to vindicate deep second-personality was proposed by Montaigne. This view claims that an agent's ends can be constituted through her relations with others in such a way that her relationships with those others always remain an ineliminable feature of her justification for pursuing those ends. Montaigne's thoughts on friendship suggest a conception of self-hood as collaboratively constructed in dialogue with one's friends. It is this role of the friend as the co-author of the self that provides the basis for this second argument for deep second-personality. Again for ease of reference, I will label this view *deep second-personality of the self*.

My goal will be to show that both theses about deep second-personality have implausible implications. Moreover, both views are in tension with perfectly plausible reductionist accounts – accounts which explain the second-personal quality of the respective sets of reasons as ultimately reducible to sets of first- and third-personal expressions. Thus, I will conclude that deep second-personality is implausible. Despite their implausibility when all things are considered, however, the exploration of these views will help to articulate an appreciation of the pervasive, albeit shallower, second-personality that characterises swathes of human practical reasoning.

---

[25] In general, I will use the expression 'reasons *for* love', to refer to reasons for loving a person, and reserve 'reasons *of* love' for the practical reasons that might arise from one's loving relations.

1. A taxonomy of second-personal reasons

Let us say that A's relation to B is second-personal if and only if A sees B as a partner in intentional interaction, such as address.[26] Thus, A's relation to B is not second-personal if A sees B as a merely potential partner in interaction, but only if A sees B as an actual partner. Intentional interaction is such that A seeing B as a partner in such interaction requires A to take B to see A in turn as such a partner. This goes some way to establishing the mutuality that is often thought to be central to what second-person relations are.[27] To be clear, on the working-definition that I am proposing, B need not *actually* see A in the way that A thinks that B does (which is also the same way that A sees B), in order for A's relation to B to be second-personal. Actual – rather than merely apparent – mutual recognition is not required for the one-directional relation of A-to-B to be second-personal, nor for the one-directional relation of B-to-A to be second-personal. But it follows from my suggested definition that for the second-person relation to hold in both directions, there must be *actual* mutual recognition. The reason that, as per my definition, I still think it worth calling A's one-directional relation to B second-personal, even if it only depends on merely apparently mutual recognition, is that from A's perspective this will be indistinguishable from the multi-directional second-person relation. This is because from A's perspective there is no difference between B's seeming to see A as a partner in intentional interaction, on the one hand, and B's actually seeing A in such a manner, on the other. And since I am interested in an analysis of the impact of second-person relations on agents' practical reasons, the relevant sort of relations must be in principle epistemically accessible to the agents whose reasons they affect. Whilst this definition of second-person relations will not find assent from every single commentator (recall the wide variety of understandings of the concept of the second-person that were canvassed in the Introduction to this dissertation), it will nonetheless provide a sound basis for understanding a broad array of the practical significances that second-persons as such can have.

There are two definitions that will help to clarify the proceeding discussion. The first, just established, is what it is for a relation to be second-personal. To reiterate, my answer is that A's relation to B is second-personal if and only if A sees B as a partner in intentional interaction. The second required definition is of a second-personal reason. Let us say that a practical reason is second-personal if and only if it is a fact, P, which speaks in favour of an agent, A, performing an action, φ, only in virtue of that agent standing in a second-personal relation, R, to some other person. Expressed more neatly:

Iff [(P favours φ-ing for A) only if R obtains], then P is a second-personal reason.

---

[26] See (Haddock 2014) and (Rödl 2014) for discussions of the second-personal nature of intentional transactions. See also (Peacocke 2014) for a sceptical take on role of mutual recognition in giving rise to distinctly second-personal thought, and (Martin 2014) for a response.

[27] Take, for example, Zahavi (2014, pp. 246-7), who takes himself to be channelling Husserl: 'the second-person perspective involves a relation between you and me, where the unique feature of relating to you as you is that you also have a second-person perspective on me, that is, you take me as your you. To that extent, there cannot be a single you: there always has to be at least two. In short, to adopt the second-person perspective is to engage in a subject–subject (you–me) relation where I am aware of the other and at the same time implicitly aware of myself in the accusative, as attended to or addressed by the other'.

With these definitions in hand, it is possible to see that there are a variety of ways that reasons could depend on second-person relations. That is to say that there are a variety of ways that the clause within the square brackets of the above expression could be realised. In the remainder of this section, I will propose a series of schematic distinctions for understanding this variety. In sketching this taxonomy, I am not endorsing the idea that all of these species of second-personal reasons really are reasons, or that they really are second-personal. Rather, the schema maps a host of second-personal reasons that *might* be thought to exist. Whether they do in fact exist is a matter that I will go on to survey in the ensuing discussion below.

The chart plots the range of different senses and extents to which a reason can be second-personal. Each box identifies a category of reason for action which are dependent on second-person relations. The chart progresses downwards by distinguishing between various subcategories within the higher-up categories, where each of the distinctions between sub-categories exhausts the conceptual space within the higher category. For example, any reason which is categorised within division 1.1, must belong in either division 1.1.1 or 1.1.2. By identifying some reasons which are second-personal in one sense, but whose normativity is also significantly grounded by further non-second-personal considerations, and by identifying contrasting sets of reasons whose normative grounding seems to involve second-person relations at every level, the chart articulates a range of *degrees* of second-personality.

2. Commands from the powerful, and advice

The first distinction that I propose as a useful one for making sense of the variety of second-personal reasons is between reasons that depend for their normativity on being communicated in a particular instance of address, and those that do not. Those that do, which will be grouped on the left-hand side of the chart in Figure 2, include: requests, commands, demands, promises, offers, invitations, apologies, issues of consent, and advice, perhaps among others. I will return later to discuss the right-hand side of the chart, where the reasons in question do not depend on particular instances of address. For now, though, it is worth delving deeper into the senses in which – and degrees to which – the reasons just listed are second-personal.

What the reasons on the left-hand side of the chart (i.e. in division 1) all share is just that they depend on a particular instance of address. Compared with the putatively irreducible second-personality of the moral demands that Darwall discusses in *SPS*, this mere dependence on address may seem to endow these reasons with a kind of second-personality that is relatively superficial. Think, particularly, of a command issued by a sous chef, Doreen, to a kitchen porter, Becky, to clean a certain surface in the kitchen. In a way, the reason Becky now has – one that was not there prior to the uttering of the instruction – is grounded in the authority that Doreen just happens to have by virtue of occupying the position of sous chef in this kitchen's chain of command. In turn, that role is one which   does   indeed

Figure 2.

Second-personal reasons

1
Dependent on a particular instance of address

2
Not dependent on any particular actual instance of address

1.1
Standing to address is itself grounded in the context of the 2nd-p relation

1.2
Standing to address grounded in non-2nd-p facts

2.1
Metaphysically dependent on 2nd-p relation

2.2
Merely epistemically dependent on 2nd-p relation

1.1.1
Grounding not reducible to commitment

1.1.2
Groudning reducible to commitment

2.1.1
Deep second-personality of the other

2.1.2
Deep second-personality of the self

create real reasons for Becky to do as she's told, but only because of some further instrumental commitment that Becky has to respecting that chain of command.[28] She may, for instance, respect this authority just as a means to getting her wages, or also as a means to producing good food for the restaurant's clientele, and / or as a means to preserving the reputation of the restaurant with which Becky identifies herself and in which she takes some pride, or whatever. The basis of Doreen's power is incidental to her second-personal relation with Becky, rather than being essential to that relation.

However, it is worth stressing that even this comparatively shallow kind of second-personality shared by everything in division 1 really is a kind of second-personality. Attending to this species of reasons may not fully capture everything that is alluring in Feuerbach's proclamation, but it does reveal *something*: namely, that to the extent that we can be rightly moved by reasons that depend on the specific addresses of others who we relate to second-personally, our agency is such that it operates in a world that we share with each other. We share our agential world not only in the sense of being in it together, but in the sense of affecting it with our intentions, which are objects of our mutual self-consciousness. That is, in the case of a command, Doreen's intention to create a reason for Becky is recognised by both as itself legitimately creating the intended reason.[29]

Suppose then that the power that Doreen has to create a reason through her address to Becky is, as suggested, grounded in facts that are not about Doreen and Becky's mutually recognising one another. In this case, Doreen's command exemplifies the second-personal reasons that are further categorised into division 1.2. The reasons in 1.2 have nothing further second-personal about them beyond the fact that they were, necessarily, brought about by an act of address. In these cases, the standing that the addressor of the reason required in order to be seen by the addressee as being sufficiently authoritative to create the reason, is not itself in any way second-personal. They remain second-personal insofar – and only insofar – as this power to give the reason was imbued in the person of the reason-giver, to be exercised through second-personal address. We might say that however that power is grounded, it is exercised by the addresser *qua* second-person.

Not all commands obviously fall into division 1.2, and not only commands fall there either. Sometimes, the power people have to create obligatory reasons for others to do things by the very expression of the intention to create such reasons – i.e. the power to make commands – might plausibly be fully justified only with reference to the relationship between the addressor and the

---

[28] By focusing on its grounding in Becky's instrumental commitment, we can understand an aspect of the second-personality of Doreen's authority that is quite separate from the second-personality that *all* legitimate authority might be thought to have. As mentioned already, Darwall thinks of all legitimate authority as irreducibly second-personal in that it is conceptually tied, he thinks, to second-personal address, valid claims, and accountability. Be that as it may, my point in this discussion is to isolate some separate, further senses in which the authority to make commands can be second-personal (or not, as the case may sometimes be).

[29] See (Enoch 2011) for a more elaborate discussion of the role of intentions in the creation of reasons by commanding.

addressee.[30] And on the other hand, some reason-giving powers other than commands are justified fully without any further reference to the second-person relation. This might well include such reasons as those addressed through offers, where the thing offered just so happens to be the addressee's to offer. Similarly, the normative grounds of many typical invitations and issues of consent may be given without any reference to second-person relations, beyond that required for the actual exercising of these normative powers. For example, if Olwen offers Josie her protractor, the mere fact that it is Olwen's protractor is all that Josie need acknowledge in order to grant Olwen the power to make the offer (felicitously). And the mere fact that it is Olwen's protractor is not a fact that adds any further second-personal quality to this reason than it has by virtue of being created in an instance of second-personal address.

One might also think that some or all reasons given through advice will fall here into division 1.2.[31] The normative structure of advising is complicated, because of its connection with trust. For my purposes, it will be sufficient here just to outline some of the possible shapes that normative structure could take so that one can see the possible ways that advice – and trust more broadly – could be second-personal. The philosophical literature about trust often focuses on the tension between trust seeming to be a rational, or cognitive attitude, on the one hand, and on the other, its being deontic. Suppose, for example, that Millie advises Gilaine to φ. If Gilaine knows that Millie is in a good position to judge whether or not she (Gilaine) should φ – if Millie is an expert in the matter, for instance – then she might well trust Millie's advice. In this sense, that is, insofar as Gilaine's trust is based on Gilaine's judgement of Millie's reliability as an epistemic source of information about the matter of whether φ-ing is a good idea, Gilaine's trust is a rational, or a cognitive attitude. It is the kind of attitude that in her canonical paper, Karen Jones (1996) would distinguish from trust proper, because Gilaine's rational trust does not yet have anything to do with Millie's dispositions to care about whether Gilaine follows her advice. Here is a kind of advising, then, which is only very minimally second-personal. It falls squarely into division 1.2.

What this contrasts with, is a species of trust that is more second-personal. It is quite possible for Gilaine's reason to follow Millie's advice to be more than merely rational, or perhaps for it to be of a different nature altogether from prudential rationality. Perhaps the paradigm instances of advising are instances in which the advisor (Millie, in this case) takes themselves to be giving a reason which invokes the relationship between advisor and advisee. Thus, Millie could feel hurt, or let down, or perhaps even angry, if Gilaine were to ignore her advice. Similarly, one can imagine Millie, not as an expert but as a friend, saying something like, 'I can't tell you why you should φ, nor how I know that you should, but you have to trust me.' The fact that Millie could conceivably say such a thing could be taken to suggest that there is a species of trust that is not just a cognitive attitude. This not-merely-

---

[30] A view close to this is advanced in (Wallace 2012).

[31] See (Faulkner 2014) for a further elaboration of this argument, that reasons of trust – including advice – are not second-personal all the way down, so to speak.

cognitive kind of trusting gives rise to reasons that are keyed to the (second-personal) relationships between the parties. If reasons of trust of this sort do exist, then rather than belonging in 1.2, they will be sorted under 1.1 in my chart.

There are numerous further complications in which one could become embroiled here. For example, one might contend that all reasons of trust are ultimately purely cognitive and rational, even if, in deciding whether to trust someone, we make appraisals of their cognitive and evaluative capacities where those appraisals themselves – unlike the reasons of trust – are deontically loaded moves in the dynamics of interpersonal relationships.[32] Accordingly, it could be that Gilaine's decision not to follow Millie's advice is only ever properly made on the basis that Millie is unlikely to be right about whether Gilaine should φ. If this is the correct way to evaluate it, then the putative reason given in the advice was only ever a reason of the sort that belongs in 1.2. But this need not rule out the possibility that in deciding to ignore Millie, and thereby treating her as untrustworthy, Gilaine offends Millie in a way that can only be understood as a transaction within their mutually-recognitive relationship. The point of this contention is that whatever reason one might have in Gilaine's position not to offend one's advisor, that could be quite separate from the reason that one's advisor gives – and takes themselves to give – in their advice.

In any case, my goal is not to resolve the matter of the nature of trust. I will settle for having illustrated the appeal of the possibility that there are species of reason-giving that are familiar through our everyday experiences of giving and receiving advice, which fall in divisions 1.2 and 1.1 on the chart. A little later on I will turn to a kind of reason which bears a close resemblance to advice, in discussing the deep second-personality of the self. The reasons we have just been considering arise as a result of actual instances of advice, or of A somehow communicating to B that B should trust A that B should φ. A nearby but distinct kind of reason might be the reason B would have to φ in virtue of the fact that A *would* advise B to φ, even when A does not actually communicate the advice. If there really are such reasons, above and beyond whatever epistemic reasons there are to think that φ-ing is a good idea, then they do seem to have a deeper kind of second-personality that is worthy of consideration, to be discussed later in this chapter.

3. Requests, promises and loving commands

Requests, promises, and loving commands are representatives of normative powers that are born out of certain kinds of relationships, such as friendship. The second-personality of these reasons runs a little deeper than that of the sorts of commands and advice discussed in the previous section, and that situates them in division 1.1 of the chart. Because whereas the reasons created in those kinds of commands and advice are second-personal in the sense that they are brought about through acts of address, their normative force is grounded in some further considerations that is not second-personal

---

[32] For elaborations of such rationalist views of trust, see (Hieronymi 2008) and (McMyler 2011).

at all, such as the power of the commander *qua* employer, or the reliability of the advisor *qua* expert. The fact that the commander has the power to, say, fire the commandee, is totally independent of any mutually recognitive relationship between the two. That is to say, Doreen the sous chef could fire Becky the kitchen porter even if Becky were to disavow Doreen's power to do so. Similarly, if Millie is an expert in a given field, and Gilaine a novice, then the fact of Millie's expert status (relative to Gilaine) is not dependent at all on any second-personal relation between them, unlike the actual addressing of a piece of advice from Millie to Gilaine.[33]

Requests, promises, and loving commands are different precisely because they are grounded in further facts arising from second-person relations, beyond just the second-personal facts that enabled the communicative acts through which these reasons came into existence. For the time being, my goal is not to give an exhaustive or systematic account of the normative structure of these three kinds of reason-giving. In Chapter Four I will return to the subject of requests to give a more thorough account of this important (because paradigmatic) kind of second-personal reason-giving. For now, it will suffice instead to illustrate the possibility of these kinds of reason-giving as dependent on this deeper sort of second-personality, whilst highlighting the differences between the three.

People entwined in personal relationships with one another typically grant one another a variety of normative powers.[34] Such powers, for example, include the power to create 'voluntary obligations' by making promises. This much has been a focus of philosophical and legal thought since at least Grotius and Pufendorf in the seventeenth century. I would like here to draw attention to the variety of such normative powers, and to the fact that they occupy an intermediate tier in the hierarchy of how deeply they depend on second-personal relations. They occupy this intermediate status because they depend on second-personal relations in a deeper way than does mere commanding or mere advising, but at the same time, unlike the forms of deep second-personality that I will come on to discuss, the normative powers at issue need not necessarily be second-personal all the way down. That is the significance of division 1.1 in the schema.

Let's suppose that when Jacqui makes a reasonable request of her friend Crystal to φ, Crystal is presented with a reason to φ that she did not have prior to Jacqui making the request. (Again, I will defend this supposition in the next chapter.) The reason thus created is second-personal not only because the request as a transaction can only be made when the addressor and addressee recognise one another in the appropriate way, but also because it is in virtue of the value that Crystal sees in Jacqui as her friend that Jacqui has the power to enact this transaction efficaciously: that is, to actually create a reason. The value for Crystal, of Jacqui, as a friend, is itself second-personal at least in the sense that it is a product of mutually recognitive interactions between Crystal and Jacqui.

---

[33] To be sure, my claim here is that Millie's expertise is not dependent on being recognised *by Gilaine*. This is compatible with the view that expertise is the kind of thing that is partly constituted by social recognition. For a discussion of this kind of role of recognition in constituting certain features of people's identities, which has a tangential relevance to the discussion at hand, see (Brandom 2007, esp. pp. 37-onwards).

[34] For clarification of the notion of a normative power, see (Raz 1999, pp. 97-106).

Friendship is the kind of thing that thoroughly depends on parties relating to one another in a second-personal way. Friendship is itself a practice and a mode of interaction in which each person must see the other as a potential partner, and seeing another person as a partner in this kind of interaction just is a way of relating to them second-personally (in accordance with my definition given earlier). Moreover, friendship is a representative of a form that second-personal relations can take in which the second-person relation itself can go some way in grounding normative powers that parties can have to intentionally create reasons for one another to act. Requesting is one such normative power.

Another such normative power, closely related to requests, is the power to create what I am calling loving commands. These commands are different from those discussed in the previous section in that the power that the commander has to give a reason is not grounded straightforwardly in anything so third-personal as, say, their capacity to have you fired, like the command of Doreen the sous chef. The power to make loving commands, by contrast, is grounded in a value that resides in the relationship between addressor and addressee. To that extent they are similar to requests. They are different from requests, though, because commands create obligatory reasons, whereas requests create reasons that are discretionary. To be clear, these terms ('command' and 'request') as they are naturally used are broader and each can be intelligibly used to refer both to reasons that are obligatory and others that are not. The point rather is that some reasons that are created intentionally though an act of address and which are grounded somehow in the relationship between addressor and addressee are obligatory, others are not. Let us stipulatively refer to the former as commands and the latter as requests. I will discuss and defend this use of the terminology at greater length in my treatment of requests in the next chapter. What matters here is that there are these two distinct kinds of reason-giving power at 1.1: requests and loving commands.

Attempts to vindicate the normative power of commanding may well stand or fall completely independently of an account putatively vindicating requests. Because however it is that our relationships can sometimes be such as to endow parties with the powers to create obligations for one another, it is not obvious that such powers entail the further power to create discretionary reasons. This may seem surprising given that the power to create discretionary reasons looks to differ from the power to make commands only by being a weaker version of the same power. But to see the more fundamental difference between the two, consider as an example the respectful if somewhat frosty relationship between two neighbours, Doris and Gladys, who for years have been the only two remaining inhabitants of a remote, rural hamlet. Doris and Gladys are both thoroughly misanthropic. They prefer for the most part to leave each other well alone, and go about their respective business in solitude. But despite all this, they grant each other the power to make certain commands – *loving* commands, in my terminology, even though there might not seem to be very much love involved.

For instance, Doris has become increasingly bothered by how unruly Gladys' front garden has grown – it directly borders Doris' own front garden, which she likes to keep in immaculate order.

Indeed, Gladys has noticed Doris scowling at the sight of her (Gladys') shabby hedge, and gangly front lawn. Doris' evident desire for Gladys to have a little pride in appearances and to do some gardening is not sufficient reason for Gladys to oblige. In fact, Gladys is rather amused by what she sees as Doris' petty concern with such trifling, cosmetic matters, and even indulges in a certain schadenfreude in being the object of Doris' irritation. Doris' mere desires, then, do not carry much practical weight for Gladys. Her loving commands, though, are another matter. When Doris after some weeks of conspicuously loud grumbling knocks on her neighbour's door and, after a respectful exchange of pleasantries, comes to the point with the line, 'Mow your lawn, it's a disgrace,' Gladys nods, and, the next day, cuts her lawn.

This utterance is a loving command, rather than any other sort of command, because Doris' power to make it is grounded in the relationship between the two of them. Their relationship is such that Doris is empowered to invoke her own interest (in this case, her desire for Gladys' garden to look a little smarter), as the basis of an obligation for Gladys to undertake a particular action. It is not that Gladys thinks that Doris is pointing out an obligation that exists independently of her command. Indeed, as we know, Gladys finds the activity of gardening quite ridiculous, and the notion of obligations to engage in it, doubly so. Neither is Gladys moved to act on Doris' order out of a fear of the consequences of disobeying. We need not think that Gladys has any such fear in mind in order for us to find it perfectly intelligible that she might treat the harmless commands of her neighbour as reason-giving, just because that is the kind of power to which their respectful if frosty relationship gives rise.[35]

On the other hand, Doris' utterance is a loving command, rather than a request, because of the kind of reactive attitudes that would be licensed in the event of the command being ignored. Following Darwall, as discussed of the previous chapter, the line between obligatory and non-obligatory reasons can be understood as the line between reasons the ignoring of which licenses blame, and reasons the ignoring of which does not license blame. Requests are discretionary, which means that a requestee's decision not to accede to a request does not automatically license the requester to blame them. Rather, such a decision may warrant other sorts of reactive attitude: the requester might legitimately feel hurt, or let down, or disappointed in the requestee (as opposed to simply being disappointed simpliciter, though they may of course feel that way too). Doris' command, as I am inviting you to imagine it, need not be thought of as a request. If Gladys were to disrespect Doris' loving command, then Doris might well be angry, and moreover, Gladys might recognise Doris' standing to blame her. Prior to the uttering of the command, Gladys might well have felt, defiantly,

---

[35] To be sure, it *could* be that sometimes people have certain kinds of normative powers to issue others with commands, where the normative force of the command for the commandee is grounded in their rightful fear of the consequences of non-compliance. Such commands, ensuing from brute power, would not have the *de jure* legitimacy that people often take themselves to have when they make commands; neither would they be *loving* commands in my sense, since their normativity is not at all derived from the value of the relationship between parties. My point in this paragraph, then, is just to point out that whilst there can be other sorts as well, loving commands are conceivable.

that no excuse was needed for her untidy garden. But once the command has been uttered, she may well acknowledge that failure to comply would require having an excuse to make to Doris.

Furthermore, as well as it being perfectly imaginable that Gladys might grant her neighbour the power to make certain reasonable loving commands, as distinct from requests, it is also imaginable that at the same time Gladys does not grant Doris any power to make requests at all. Suppose that Doris were to attempt to make a request, and that she successfully communicated to Gladys that she was only asking her to mow the lawn, and that if Gladys decided not to, then she (Doris) would understand, and would not blame her for it. That is, we are supposing that Doris communicates to Gladys exactly what kind of reason she takes herself to be giving, not a command but a request, where the matter of whether to accede really is a matter of Gladys' own discretion, and the worst reaction that Doris would be entitled to in the event of Gladys ignoring the request would be to feel hurt, or let down, or disappointed in Gladys. If Doris presented all this to Gladys, we can quite easily imagine Gladys responding with a blank stare, and being totally unmoved by the request. That is, just because we are imagining their relationship to be such that they grant each other the power to create obligatory reasons through commands, it does not follow that we must imagine them to grant each other the further power to create discretionary reasons through requests. Indeed, granting that Doris can make efficacious commands, we can still imagine Gladys finding her attempted request to be completely ridiculous, presuming as it does that Gladys cares a jot about Doris having hurt feelings, or being let down, or being disappointed in Gladys.

The point of this example has been to illustrate the difference between the normative structures of requests, and loving commands. Part of the purpose of so doing was to show that there are a variety of distinct kinds of reason that cohabit the conceptual space in division 1.1. This is a component part of my overall goal in this discussion, which is to illuminate the variety of kinds of reasons that might be called second-personal reasons, and the variety of senses in which and extents to which they are second-personal. Requests and loving commands are second-personal in a similar sense, and to a similar extent. As well as depending on a second-personal relation in order to be created through an act of address, they are normative powers that draw their force from the second-personal relationships, such as friendship, between addressor and addressee. This finds them both in division 1.1.1

The third representative of this class of reasons at 1.1 is promises. These can be quite a lot like loving commands in that they are normative powers to create obligatory reasons through acts of address. However, there is a salient species of promises which help to illuminate the final difference between senses and degrees of second-personality that I want to articulate on the left-hand half of the chart: that is, the distinction between 1.1.1, and 1.1.2. The species of promises in question are such that the normative force of the ensuing obligation derives not from any valuing attitude that the promisor places in the person of the promisee, but rather from a mere *commitment* that the promisor has undertaken, and the related *rights* of the promisee. As I will show presently, such promises are in

a way less second-personal than others, and less second-personal than the requests and loving commands that have been central to my discussion in this section.

The kind of second-personal reasons that are grouped at 1.1.2 are exemplified by the reasons generated in legally binding contracts. Or at least, such reasons are exemplary of this division insofar as the laws that govern such contracts are themselves just laws, and thus the reasons are real normative reasons. Suppose that Tess hires Diane, a private eye, to tail her (Tess's) daughter for a month. They sign a contract agreeing that Tess will pay half the fee up front at the beginning of the month, and the other half at the end of the month in exchange for Diane's report on Tess's daughter's movements. Contracts such as these are a form of promising. Tess promises to pay; Diane promises to spy. Each is now constrained by obligations that arose from the transaction between them that could only be transacted second-personally, i.e. with each recognising the other as a partner in intentional interaction. This example just so happens to involve two promises, and thus has a balanced nature to it, but it needn't. Diane could just as well have promised – by signing a contract, for example – to undertake the investigative work even without the countervailing promise of payment. She could undertake that commitment for whatever reason of her own she might have. The point is that it is the very act of promising – not the prospect of remuneration – which generates her subsequent obligation.

Promises such as Diane's differ strikingly from the sort of promise that one can make in a close personal relationship. Consider Martha's promise to her dearest friend, June, that she (Martha) would not join the rebel militia. The reason that June insists that Martha promise not to join the militia is Martha's very own self-interest, for which she herself has too little regard (or at least so it seems to June). As a result of the promise, Martha has an obligation to June that she did not have prior to the promise, not to join the militia. In the context of loving friendships like June and Martha's, subjects may grant others the normative power to hold them accountable to commitments that they make in their own interest, or in the interest of the other, or in that of the friendship itself. To be sure, there are complexities clouding the normative structure of promises like Martha's.[36] What is quite clear, though, is that the basis and normative power of Martha's command is not only different from that made by Diane the private eye. Martha's is also *more* second-personal. Whereas Diane could have made and been normatively bound by her promise to Tess irrespective of how she felt towards Tess – whether she felt nothing but spite and loathing towards her, or had no thoughts towards her whatsoever – Martha's promise is different. It can only be efficaciously made as the kind of promise that it is in virtue of a relationship between herself and June which is more than a mere commitment. It is a robust friendship. As such, the promise is partly grounded in the value that each acknowledges the other to acknowledge each as having, not only as people, but each as a friend to the other. Evidently, this kind of promise is more deeply second-personal than reasons grounded in mere commitments. So too are the sorts of requests and loving commands already discussed which draw from the normativity of the

---

[36] For a discussion of some of these complexities see (Owens 2006).

valuing relationships between parties. Such is the difference between divisions 1.1.1 and 1.1.2 on the chart.

So concludes my discussion of the varieties of second-personal reason-giving on the left-hand side of the chart. I have shown that there are many ways that people can alter the normative landscape through acts of second-personal address. Contra Darwall, making a moral demand is not one of those ways, as I argued in the previous chapter. But among the normative moves that can be made through address, some are more second-personal than others. Some invoke powers that a party has incidentally to the relationship between addressor and addressee (like the power to lend a protractor, or fire a member of staff, or give a good piece of advice). Others invoke relationships between parties that consist in a single, bare commitment, as exemplified by legally binding contracts. Still others, though, invoke the value that people have for one another in personal relationships which value itself is dependent on layers and histories of second-personal, mutual recognition.

4. The relationality of personal value

All of the reasons on the left-hand side of the chart satisfied my definition of a reason being second-personal at least because they all depended for their normativity on being addressed, which is a transaction that itself requires addressor and addressee to be related second-personally. The reasons on the right-hand side of the chart, however, do not depend on any such address. As such, the senses in which these reasons are second-personal must be different. In this section I will animate the plausibility of there being reasons stemming from the value of particular people, epistemic access to which can only be attained via second-personal relations with those people. Such reasons comprise division 2.2 in the chart. This possibility – that second-person relations can play an essential epistemic role in reasons for love (perhaps among others) – will be important later on because it offers an explanation of the practical significance of the second-person relation for reasons for love, which functions as a rival to that explanation advanced by the view that I will call deep second-personality of the other.

To show the possibility of there being some reasons for loving people that are in any sense second-personal reasons, a little must first be said to clarify what reasons for love might be. Within philosophical debates about the nature and normativity of love, there are some (most prominently Harry Frankfurt (2004)), who deny that love is the kind of thing for which reasons could be given. That is only one view though, and in what follows I will be assuming, contra Frankfurt, that love is a rational attitude for which reasons could, and indeed should (in any instance of loving) be available.[37] Among those who agree that love is rational in this sense, there are broadly three views about the kind of thing that reasons for love are. One is that all of the things about an individual that make them worthy of love are generic traits that the individual shares with every other individual: their bare

---

[37] I will revisit this assumption and give a defence of my rationalist stance on reasons for love in Chapter Five.

humanity, or their personhood, or some such. This view advances the thought that simply being a person is all there is to being worthy of love, and that there is nothing rational at all about loving some people and not others on the basis of things that make such people different from such others. I will refer to such views as *Bare Personhood views*.[38]

Against the Bare Personhood views are those who think that when an individual is worthy of love, that is at least partly in virtue of qualities they bear that mark them apart from others, that is, in virtue of the things that make them who they are (as opposed to being someone else, or say, a less-good version of themselves). Let's call all such views *Qualities views*.[39] Qualities views are committed to the idea that because people have different individual qualities from one another, and because it is those very individual qualities (perhaps in conjunction with generic qualities like personhood) that render love appropriate, some people may be more worthy of love than others.

Within the Qualities views, there are two positions that are worth treating as distinct for my purposes, for reasons that will become clear. One position is that the individual qualities that warrant loving their bearer are moral qualities.[40] Let's call this the *Moral Qualities view*. It is, on the Moral Qualities view, a person's kindness, their honesty, their courage, and so on, that could speak in favour of loving them. Other qualities that they bear might be indicative of their moral qualities – the way that a person's good humour might be an expression of their patience and compassion, for example. But insofar as a person's wit, say, is not indicative of any distinctively moral qualities of theirs, it is not something that rationally favours loving them, on this view.

Naturally enough, what the Moral Qualities view contrasts with is a view – call it the *Charming Qualities view* – that the individual qualities of a person that can speak in favour of loving them are not only moral qualities. Thus, non-moral charms such as a person's wicked wit, or their passionate nature, or their creative talent, can also make it appropriate to love them. Whereas some such non-moral qualities – like the ones just listed – might sound like the kind of things that are clearly desirable in a person and which intuitively make a person worthy of love, the charming qualities view can also incorporate the way that in some loving relationships it is apparently mundane qualities that seem to be the basis of the love. For instance, it is not uncommon to hear of lovers claiming to love their beloved's defects and imperfections. The Charming Qualities view has scope to argue that various characteristics that people find to be endearing or loveable despite their *prima facie* lack of positive value, might play a role in constituting a person's charm when all things are considered.

---

[38] For proponents of bare personhood views, see (Velleman 1999; Setiya 2014). The *locus classicus* for this idea is Kant's *Groundwork* (1997, e.g. 4:425)

[39] This terminology follows (Jollimore 2011).

[40] A spectre is raised by the mere mention of a distinction between qualities, and their bearer. Some people have worried that Qualities views are doomed because they will always wrongly characterise love as directed towards a person's qualities, rather than to the individual themselves, the bearer of the qualities. I am persuaded by Christopher Grau's (2004, pp. 120-1) line of response to this problem, which is to note that it is not a problem that should bother philosophy of love any more than it should bother any other area of philosophy where 'bearers' are thought to have properties. This I think is enough for the Qualities view to have banished the spectre of metaphysical confusion.

Thus, between the Bare Personhood view, the Moral Qualities view, and the Charming Qualities view, there are a range of perspectives on the nature of reasons for love which are live contenders in the contemporary philosophical debates about love. In the remainder of this section I will seek to show how second-person relations might be thought to play a necessary epistemic role in revealing to a subject the love-worthiness of another person on any of these three accounts of what such love-worthiness consists in. I will do so by presenting this epistemic function of second-person relations when the Charming Qualities view is assumed, and then by arguing that the plausibility of the Bare Personhood and Moral Qualities views depend on those views being similar in relevant ways to the Charming Qualities view.

A further preliminary clarifying remark to make before proceeding is about the kind of love that is under discussion. There are, after all, normally thought to be wide differences between a kind of neighbourly love that is universal, and personal love with those with whom one is very close.[41] And moreover, there are often thought to be further significant differences between mature romantic love, immature romantic love, friendship, parental love, filial love, brotherly and sisterly love, and so on. The purpose of my discussion in this chapter in general, and this section in particular – besides disambiguating a number of senses in which reasons can be second-personal – is to offer first-blush defences of the existence of a number of those different kinds of second-personal reasons. As such, if there are *any* species of love which are rational and whose reasons might be epistemically accessible only within second-person relations, then those reasons are in that sense second-personal and they belong in my schema in division 2.2. However, it is worth noting that some of the literature that I have already referred to in connection with the three schools of thought regarding reasons for love focuses heavily on reasons for loving people in a (mature) romantic sense. There is a presupposition underscoring some of these philosophical debates that romantic love – particularly its full-blooded, mature form – is so importantly different from other species of love that the discussion of its rational structure must be quite independent from theirs. This is associated with the further thought that romantic love is characterised by the way one's beloved fills one's heart completely and singularly so that for someone who is truly in love there could not possibly be room for anyone other than one's beloved. This peculiar phenomenon is thought to warrant a distinctive explanation, and whatever there is to say about the rationality of, say, our love for our friends, it will be something altogether different. For what it is worth, I am generally sceptical of this entire line of thinking, which to my mind dramatically exaggerates the difference between romantic and non-romantic love. Though it is beyond my present scope to defend this contention here, I suspect that the only differences that could stably distinguish romantic from non-romantic love are the latter's lack of desire for or realisation of sexual activity between lover and beloved. At the very least, let it be clear that the kinds of second-personal

---

[41] In fact, I will take up this thought in Chapter Five and argue that agape – or neighbourly love – is much closer in nature to personal love than is usually assumed.

epistemic access to reasons for love under discussion here should not be thought of as restricted to romantic love, but also to include other personal loves such as friendship.[42]

With this breadth in mind, let us suppose for a moment that the Charming Qualities view is the correct one: that the non-moral charms that a person might have (which may distinguish them from others) are among the kind of facts that could speak in favour of loving them. The relevant question that must be addressed to appreciate the second-personality of this kind of reasons for love is how one comes to see that array of non-moral features, and how one comes to appreciate such features as charms. The crucial observation that I want to submit about loving someone in virtue of the peculiarities that make them who they are including perhaps their limitations, ticks, vices or imperfections, is that it may often be very difficult or even impossible to render them explicit as one's reasons for love.

This fact that a person's non-moral charms might well be difficult to articulate warrants some further exploration. To be clear, there are some non-moral charms that might come readily to mind when one considers one's reasons for love. Suppose that Delyth loves Heulwen, and when she asks herself why she loves Heulwen, Delyth immediately thinks of Heulwen's almost impulsive spontaneity, and the mirth on her face with she flashes a trademark look at Delyth. Such qualities as these might be easy enough for Delyth to explain as reasons for loving Heulwen: they are straightforwardly positive qualities (or indicators thereof), they are consciously accessible and conceptually articulable. However, Heulwen is also wont to be a little short-tempered, and more than a little disorganised. From Delyth's point of view, it is not that she loves Heulwen in spite of these short-comings. They are not negative entries in the ledger of her all-things-considered love-worthiness which just happen to be outweighed by her charms. Rather, for Delyth these apparent short-comings *are* charms: she loves Heulwen *for* being occasionally curt, and persistently chaotic. It is difficult, if not impossible, though, for Delyth to explain why such apparent short-comings are in fact charms. This inexplicability represents one aspect of the inarticulability of reasons for love on the Charming Qualities view.

A further aspect of that phenomenon is that some reasons for love might be not just difficult to explain as reasons for love, but difficult to articulate conceptually at all. If pressed to give a full account of why she loves Heulwen, Delyth may find to her frustration that there is just *something about her*, that draws Delyth's affection. This inarticulability might not be surprising when one considers that one's acquaintance with someone one loves might be mediated by a range of senses, and one's impression of a person as attractive, or loveable, might itself be made in several different sense modalities. Thus, the firm press of Heulwen's hand, or her smell, or the timbre of her voice may

---

[42] The notion of *personal loves* is supposed to include loving relations in the context of personal relationships such as romantic relationships, friendships, and family, but it is supposed to exclude any loving relations that are impervious to personal relationships, such as neighbour love. That distinction is adequate for the purpose of the present discussion. However, in Chapter Five, I will argue that neighbour love too ought to be understood as a kind of personal love.

all strike Delyth as parts of the explanation of the aptness of her love of Heulwen. Doubtless we do, as a matter of routine, appreciate aesthetic qualities in others that are expressed in ways that are at once too fine-grained, and too far from linguistic thought to be conceptually articulated.

To return to the question at hand, it has now become one of how it is that one can gain epistemic access to the love-worthy charms of another, where some such charms may be the sort of inarticulable qualities in virtue of which Delyth loves Heulwen. When it is made clear how nebulous reasons for love can seem, it is also clear that epistemic access to such reasons may well depend on *personal acquaintance* with the object of the love. That is, it is plausibly impossible to come to know about the aptness of loving Heulwen by any means other than becoming personally acquainted with Heulwen's love-worthiness oneself. Mere second-hand descriptions of Heulwen could in principle never fully convey the *something about her* that is so salient to Delyth.

Recognising the essential epistemic role of personal acquaintance in coming to appreciate reasons for love is a significant step towards recognising the second-personality of such reasons for love on the Charming Qualities view. But that second-personality has not yet been established because personal acquaintance does not necessarily imply a second-person relation. That is, it might be possible to become personally acquainted in the relevant sense with the inarticulable love-worthiness of another without treating them, or being treated by them, as a partner in intentional interaction, and so without relating to them, or being related to, by them, second-personally. Perhaps one might come to see a person's lovable qualities in an embodied interpersonal encounter that lacks mutual recognition – e.g. one where they do not know you're there, and you know it. Perhaps one might be struck by the love-worthiness of a particular actor, who one has seen on screen many times, only when one sees them perform on stage, for instance. If this were indeed possible, and the reasons for love on the Charming Qualities view were in principle all epistemically accessible via mere personal acquaintance with such reasons, then those reasons would not be second-personal, at least not in the sense under discussion here and captured in the chart's division 2.2. Indeed, 2.2 would be an empty division.

What is further required in order for such reasons for love, on the Charming Qualities view, to be second-personal in an epistemic sense, is for second-person relations themselves to be necessary enabling conditions for the cognition of inarticulable love-worthy qualities. Whilst this is a further step than thinking that mere personal acquaintance could alone enable the relevant cognition, it can also claim some plausibility. As will be discussed later in connection with deep second-personality of the other, some theorists have found all love-worthy traits to be relational in the sense that a trait comes to *be love-worthy* only in the context of its expression by one person to another when the two are bound in a relationship (one that is, amongst other things, second-personal). The idea that such traits can only be properly *appreciated* as love-worthy in the context of the relationships within which they are expressed is a substantially more modest claim. Think again of Delyth and Heulwen. I earlier mentioned, as an example of the things Delyth loves about Heulwen, Heulwen's trademark mirthful

look. Think here of a kind of laughing, knowing look: a non-verbal expression which laughs about something and at the very same time acknowledges that you and I are laughing about it together. It strikes me as eminently plausible that interactions like these can play a role in someone coming to love someone else, and that the role they play could be that of helping the would-be lover come to appreciate (though not necessarily in any conscious or intellectualised way) the aptness of loving this particular other. Indeed, granting that such interactions as these knowing looks can play such a role in seeing others' love-worthy qualities, it is hard to imagine how any other mechanism at all could bring those same qualities to light. From Delyth's perspective, there is a stark difference between spotting Heulwen share a mirthful look with someone else, on the one hand, and on the other, for Heulwen to share a mirthful look with Delyth herself. The former might help Delyth to imagine Heulwen's charms, the latter is itself the moment of those charms being made manifest for Delyth in such a way that she can become directly acquainted with them.

Strictly speaking, one example would suffice to show that there can be reasons for love for which second-person relations are necessary epistemic enabling conditions. But knowing looks are not a rogue, one-off example. In fact, there are many similar ways in which a person's inarticulable love-worthiness might be expressed directly to and *for* a would-be lover only in the context of second-personal interaction between the two. Think of virtues like compassion, patience, loyalty and honesty. Delyth might have heard all about Heulwen's profound honesty, and maybe those mere reports might have endeared Delyth to her. But such endearment is nothing next to the effect of Heulwen's being profoundly honest with Delyth herself. There are thus a range of reasons for loving others – paradigmatically the inarticulable 'something about' someone – which are epistemically accessible only in the context of second-person relations, at least on the Charming Qualities view about reasons for love.

Before concluding the section, I would like to turn briefly to the Moral Qualities and Bare Personhood views in the hope of being able to repeat the previous sentence but without its 'at least' clause. The Charming Qualities view was helpful in drawing out the nebulous nature of some of the things about others that we sometimes take to warrant loving them. Qualities like their smell, or the distinctive firmness of their touch, might possibly figure in justifying love on the Charming Qualities view. Such qualities are unlikely to be able to figure in such justifications on the Moral Qualities view, where it would take some story to be told explaining how a person's firmness of touch, or even their smell, reflects some moral virtue that they possess. This seems far-fetched. Proponents of the Moral Qualities view, in cases where no such story connecting a trait to a moral virtue is available, deny that the trait in question could legitimately figure in a judgment of whether love is apt. But my discussion of second-personal relations providing epistemic access to a person's inarticulable loveable 'something about them' has already touched on the idea that moral virtues could comprise part of that 'something about them'. I gave the example of Heulwen's honesty making an immeasurably stronger impression on Delyth as a reason to love her when it is Delyth to whom Heulwen is being honest.

Many other moral virtues, also mentioned above, seem to be the kinds of qualities that a person bears *through* acting certain ways with and towards others in second-personal interaction. And when such virtues are expressed through such interaction, they are presented to the partners to such interaction in a way that they could not be presented in any other, indirect, non-second-personal manner. Thus, on the Moral Qualities view as well as the Charming Qualities view, there are epistemically second-personal reasons for love occupying division 2.2.

What about the Bare Personhood view? Is there something of a person's bare personhood that is bundled up in that 'something about them' which is felt to be salient by those that love them, and which is revealed only through second-personal interaction? There are two ways that a proponent of the Bare Personhood view could respond. One way would be to answer the last question negatively, and to claim that the generic qualities that all people share which justify loving them are such that their instantiation in any particular individual can be known perfectly well without second-personal interaction with that individual. Identifying someone as a person, which can presumably be done without interacting second-personally with them, already provides one with all the grounds that one could have to love that person. So says this first strand of response from the Bare Personhood view.

Such a response has the consequence of making the Bare Personhood view sound quite far removed from what we talk about when we talk about love. That is, denying that an individual's distinct *way* of manifesting *their* personhood is irrelevant to the reason their personhood provides to love them is a strikingly cold and impersonal way to conceive of the nature of love. A defender of the Bare Personhood view may have strong enough independent reasons for endorsing that view as to warrant taking on the explanatory cost of describing the phenomena in a way that seems far removed from real life. But to be sure, that is an explanatory cost.

Alternatively, the other strategy that a Bare Personhood theorist could pursue in response would be to affirm the role of an individual's distinct way of manifesting their personhood as partly grounding the appropriateness of loving them. This coheres with the way that some Bare Personhood theorists, in opposition to the Qualities view in general, will insist that it is not a person's qualities but their *themness* that is the object of love (and, they add, rationally so).[43] If the Bare Personhood theorist takes this latter line of response then they may find that for the very same reason that initially motivated the Charming Qualities view, some reasons for love are epistemically accessible only within second-personal relationships. It may be that a person's 'themness', so to speak, is itself glimpsed only via the matrix of their particularities which are conveyed in second-personal interaction between lover and beloved.

In this section I have argued that on any of the three most popular ways of construing reasons for love, it is plausible that some such reasons can only be appreciated via second-person relations. This is an aspect of the practical significance of the second-person relation. Moreover, it is a way of

---

[43] See (Kreft n.d.) for such a line of thought.

explaining the *relationality* of a person's love-worthiness. Namely: A is worthy of B's love only in the context of A's (second-personal) relation to B, since A's worthiness for B's love can only become fully apparent to B in the context of that relation. This explanation of the relationality of a person's love-worthiness establishes a thorny rival for the next view that I will discuss in this chapter, the view I am calling the deep second-personality of the other. That view itself tries to account for the relationality of personal love-worthiness by positing that second-personal relations do not reveal, but in some way *constitute* a person's love-worthiness. The very existence of a plausible rival explanation of the same phenomenon will, I argue, cast this latter view into doubt.

5. Deep second-personality of the other

I have discussed one sense in which reasons that are not borne through second-personal address can be second-personal nonetheless – that was an epistemic sense. A stronger sense in which such reasons might be thought to be second-personal is a metaphysical sense: where reasons are at least partly *grounded* in second-person relations, rather than merely revealed therein. Such reasons occupy division 2.1 in my schema. One inhabitant of this division is straightforward: namely, reasons to act lovingly towards someone with whom one is entwined in a mutually recognitive loving relationship of some kind (again, using the term love broadly enough to encompass friendship). In my discussion of the left-hand side of the schema I considered normative powers like requests and loving commands that might arise from certain kinds of relationships. Given that relationships can give rise to such normative powers, they can certainly also give rise to special reasons to, say, act in the interest of the other, even when no express request or command is made (Wallace 2012).

However, there may be a limit to the extent to which such reasons refer to second-person relations in their grounding. If you have a reason to act in the interest of another in virtue of your relationship with that other, then the reason is grounded in the relationship. But if you have some further normative story to justify being in the relationship – perhaps in terms of the worthiness of the other for the kind of loving commitment that you have bestowed upon them – and if that further rationale does not itself invoke the relationship, then your justification for acting in the other's interest will ultimately be grounded in these non-second-personal facts (such as the other's worthiness for love). What this contrasts with, is a view that says that reasons for love might be second-personal all the way down, which is the view that I call deep second-personality of the other.

There is a thought that has been advanced in some of the recent literature about reasons for love to the effect that a beloved's love-worthiness may be irreducibly connected to their relationship with a particular other. I take it that something like this view is implicit in Kolodny (2003) and Rorty (1987). But it receives more explicit expression in Abramson & Leite (2011). They say (p. 681) that a relationship 'establishes the context within which the beloved's laudable character can become an appropriate ground for love', and, 'the relational context in which the traits are expressed, and their

expression in that context […] provide the background against which the good traits *become* an agent-relative reason for love' (my emphasis). Abramson and Leite thus claim that the relationship between A and B (which is, amongst other things, second-personal) does not just provide epistemic access to B's love-worthiness, but forms part of the grounding of B's love-worthiness. A's justification for loving B, on this view, is second-personal all the way down.

There are two objections that I would like to put forward to Abramson and Leite's view – that is, to the deep second-personality of the other. The first I have already anticipated in the previous section. Abramson and Leite's argument for their view trades on intuitions about how some forms of love are only appropriately felt towards B by those, like A, who are actually in loving relationships with B. They rightly note that B's worthiness for love is in some sense *relational* to those with whom B is entwined in close personal relationships. Abramson and Leite take this observation to speak in favour of their version of deep second-personality of the other. But this observation does not in fact speak in favour of the view that the relationship with B forms a constitutive part of the justification for loving B any more than it speaks in favour of the rival view that *only* being in a loving relationship with B can grant one epistemic access to B's love-worthiness (i.e. B's aptness as the object a particular kind of, in this case romantic, loving attitude).

The account offered in the previous section regarding the epistemic second-personality of reasons for love also offered an explanation of the relationality of a person's love-worthiness. That epistemological account thus deprives Abramson and Leite's proposals of the support they might otherwise draw from the intuitive attraction of positing that relationality of personal love-worthiness. This is one objection.

The other objection that I would like to put forward pertains to the peculiar nature of reasons for love on their account, or indeed any version of deep second-personality of the other. The problem here is that if the reasons for love are second-personal all the way down, then they do not bottom out in irreducibly normative concepts, such as goodness, or fittingness, or ought. And if the justification for loving someone does not bottom out in irreducibly normative concepts then its status as a justification is somewhat suspect. It is mysterious to posit that the mere fact that A and B are related second-personally itself counts as a reason for A to love B. How could a mere descriptive fact of this sort provide an adequate answer to the normative question of why A should love B? This line of thinking seems fairly obviously to commit the naturalistic fallacy.

Let's call the fact that A and B are engaged in a second-personal relation, R. To be clear, Abramson and Leite (and proponents of deep second-personality of the other in general) need not, and indeed do not, think of R as what Jonathan Dancy (2004, p. 38) would call a favouring reason for A to love B. That is, Abramson and Leite do not think that other things being equal, R by itself speaks in favour of A loving B. Rather, they think that R is an *enabling* reason (also Dancy's term, 2004, p. 38.). Enabling reasons are conditions the obtaining of which allow other conditions to speak in favour of some action, for some agent. Thus, Abramson and Leite think that it is B's 'laudable character' that

speaks in favour of loving B. R is just a condition the enabling of which enables that laudable character to count as a reason for loving B. Enabling reasons can be perfectly natural, descriptive facts, and indeed, there are likely many enabling reasons implicitly facilitating B's laudable character speaking in favour of loving B. A relevant example might be the fact that B's character, and their virtues, are *compatible* with A's character. Perhaps A is typically insecure about whether she herself is a good person, and whilst B is very evidently a good person, her humility is not her most prominent trait. One can imagine B being morally excellent, all things considered, but still not compatible with A. If this were the case then B's laudable character might have been undermined as a reason in favour of A loving B. Thus, the fact of their compatibility is an enabling reason for B's laudable character to favour A loving B. This is the kind of reason that R is too, according to Abramson and Leite.

However, the mysteriousness of R's normative role is not dispelled, even when it is clarified that R is supposed to be an enabling rather than a favouring reason. It is quite clear why A's and B's compatibility should be an enabling reason, because compatibility is a normative notion; it is a good thing (albeit, as an enabler, not the sort of good thing that would speak in favour of loving someone in the absence of any other reason to do so). R, on the other hand – the fact that A and B are *already* party to some form of second-person relation – is not obviously a normative notion at all.

To summarise the point of this section, it has been to mount a critique of the idea of deep second-personality of the other, the view that reasons for love might be second-personal all the way down. Abramson and Leite advance such a view explicitly, but perhaps in an implicit form it is much more widespread. Against that view, I have suggested that there is no normative role – neither favouring nor enabling – to be played by the fact of any two people being in a second-person relation with regard to their loving each other. Such second-person relations can be of rational importance in other ways: by making it possible for people to gain epistemic access to one another's charms, and indeed, to their compatibility. But these rational roles are all ultimately such that the reasons for love are reducible to facts that do not depend in any way on distinctly second-personal relations.

6. Deep second-personality of the self

In this section I will discuss the idea that relationships with particular others can contribute to one's self-identity in such a way that those relationships form ineliminable parts of one's justification for certain actions. Many writers – including phenomenologists like Merleau-Ponty, psychoanalysts like Carl Jung, communitarian critics of liberal 'atomism' like Charles Taylor and Alasdair MacIntyre, and feminists like bel hooks – have noted that our social relations play an important role in the way we as individuals construct our own identities. Through interaction with others we are inculcated into practices and traditions which offer us ways of seeing the world that presents the value of certain ends, as well as suggesting ways to realise them. It is also through our relations with others, typically, that virtues and vices are instilled in us as individuals. But these sorts of ways that social relations

contribute to the formation of an individual's selfhood can all be plausibly reduced to causal and epistemic contributions, not normative ones. That is, it is clear from these examples that other people (partially) *cause* us to become the people that we are, and that they help us to *see* the value of certain paths of action. But such examples do not suggest that interpersonal relationships themselves ever form residing parts of our reasons for *being* certain ways, i.e. for having certain self-identities which ground our courses of action.

However, one thinker who has suggested a more enduringly normative role that certain others might play in the construction of the self, is Montaigne. As such, he is a proponent of the view that I am labelling deep second-personality of the self, which occupies 2.1.2 on my schema. In his essay 'On Friendship', Montaigne recounts a story told by Cicero about a man who, on my reading, advances a view that I will call Collective Reasoning. The basic idea here is that a possible normative power friends might have would be to participate collaboratively with the subject in the formation of the subject's axiological commitments. Consider this passage:

> When the Roman consuls had condemned Caius Gracchus, and were proceeding against all those who had been in his confidence, Laelius, in their presence, inquired of Caius Blossius, who was his chief friend, how much he would have been willing to do for him. His reply was, 'Everything.' 'How, everything?' demanded Laelius. 'And what if he had ordered you to set fire to our temples?' 'He would never have told me to do that,' answered Blossius. 'But if he had,' Laelius insisted. 'Then I should have obeyed him,' said he […] [T]hose who condemn his answer as seditious have no proper understanding of the problem, and fail to take into account the fact that he held Gracchus' will in the palm of his hand, since he both influenced it and knew it. They were friends before they were citizens, friends to one another before they were either friends or enemies to their country, or friends to ambition and revolt. Having absolutely given themselves up to one another, each had absolute control over the reins of the other's inclinations. Therefore, supposing this team to have been guided by virtue and governed by reason – and it could not have been driven otherwise – Blossius' answer was as it should have been. If their actions did not fit together, they were, according to my measure, neither friends to one another nor to themselves. (Montaigne 1993, p. 98)

The relationship between Caiuses Gracchus and Blossius is offered by Montaigne as in relevant respects exemplary of the best kind of friendship. In such ideal friendship, Montaigne goes on, the minds of the friends 'mix and blend one into the other in so perfect a union that the seam which has joined them is effaced and disappears' (1993, p. 97). Some of Montaigne's imagery is hyperbolic. Surely in any friendship, ideal or otherwise, the distinction between one mind and the other will always remain apparent and their perspectives – no matter how similar – will always be different. Trivially, even if individuals share profoundly many beliefs and values, they have different pasts, different prospective futures, subtly different physical and cognitive capacities, and so on, which all figure in various ways in determining parts of their beliefs and values. Nonetheless, there is an idea worth taking seriously, I suggest, in Montaigne's conception of two wills which are controlled together, collaboratively, at the same time as being 'guided by virtue and governed by reason'.

Montaigne implements this conception of friendship's role in determining one's will, through the apparently contradictory avowals that Blossius makes under Laelius' inquisition. Blossius claims that Gracchus would never order him to burn down the city's temples, thereby implying that he

(Blossius) robustly believes that such a heinous act would not be permissible under any circumstances. And yet he further claims that if Gracchus were to order him to burn them down, then he would do it. This runs directly against the implication of his first claim. For, now Blossius implies that there could be circumstances under which it would be permissible to burn down the temples (in the rational or moral sense of permissibility, if not the legal or political sense): namely, such arson would be permissible if Gracchus were to order it. One way to interpret Blossius – and Montaigne – here is as advancing the following thesis.

Collective Reasoning: For some pair of friends, the fact that they collectively decide to φ can itself count as a reason for either to φ.

This thesis provides the resources to understand Blossius' claims. Moreover, on this view, Blossius is committed to the deep second-personality of the self. Whilst neither Blossius nor Gracchus have reason to burn down the temples, they place more trust in their relationship as a guide to right action than they place in any of their most deep-seated values, including their civic convictions. There is one way of understanding that trust Blossius has in their friendship according to which the friendship is nothing more than an epistemic guide – a heuristic – for Blossius to identify which the right course of action is. If Montaigne's conception of friendship involves friends only trusting one another in this sense, then whilst the friendships may be profoundly interpersonal means to arrive at one's axiological convictions, they are ultimately just that: epistemic means to reasons, rather than normative grounds. Collective Reasoning, however, posits a role for friendship that is genuinely normative, and not merely epistemic. It offers the following interpretation of Montaigne: that Gracchus' hypothetical command indicates that Gracchus and Blossius, whose wills are entwined, would have collectively judged that burning down the temples is the right thing to do; and that if they collectively so judge it, then it is the right thing to do.

As I have said, I think Collective Reasoning is worth taking seriously. It has some *prima facie* appeal, given a certain philosophical backdrop (albeit one that could not without anachronism be ascribed to Montaigne). I will very briefly sketch that backdrop, including an understanding of cooperative decision-making and axiological discretion that is implied by Collective Reasoning. Then, however, I will argue that even granting this backdrop, the appeal of Collective Reasoning, and so of the deep second-personality of the self, is illusory.

There are two parts to that backdrop against which Collective Reasoning can be made intelligible. One is to do with the nature of reasons, the other with the nature of collective decision-making. First, then, let's consider the presuppositions about reasons that underpin the thesis. If one were to understand the matter of what one ought to do in any and every situation to be fully determined by reasons that lie outside any decision-making of the agent, then there could be no scope for agential decisions to influence what one ought to do, and therefore similarly no role for

collectively made agential decisions either. That is to say, if a thorough going externalism about all practical reasoning were to be true, then Collective Reasoning must be false. The viability of the thesis thus depends on the assumption that at least some of the time, what one ought to do can be partly constituted by one's own chosen values and ends. To be clear, such an internalist picture of practical reasoning need not be committed to any further contentious views about the dependence of values or ends on subjective inclination. One could, for example, think that there is a multiplicity of worthy ends available to agents, at least some of the time: too many for agents to pursue. And even though the worthiness of each of these ends might be objective – in the sense of being independent from any given agent's inclinations or preferences – it could still be rationally underdetermined which of the multiplicity an agent ought to choose and commit to. Thus, there is a plausible, though not totally uncontentious, understanding of the nature of practical reasoning according to which there is some scope for agential choice in determining what one ought to do. This scope for agential individual choice may also include space in legitimate practical deliberation for Collective Reasoning.

The second background presupposition of the thesis, then, must establish that some collective decisions can be such that the fact that it was *we* that decided to φ itself is part of my reason to φ. This is not a given on every plausible sense of collective decision. A counsel of experts may gather to advise the government on a particular policy, where it is important for them to reach a clear collective recommendation. Here, though, whatever legitimate role their collective decision can play in guiding the policy decision, it will be a merely epistemic role, pointing towards the best available option. Ultimately, the fact itself that the advisory expert committee recommends that the government φ is not a reason to φ, but, at best, an indication of some other reason to φ.

Not all collective decisions are just advisory, though. There has been an expansion of interest among philosophers in recent years in the topic of shared, and group agency. There are those, such as Kirk Ludwig (2016), who think that collective agents are reducible to sets of individual agents, and thus that reasons for individuals to act need make no reference to collective decisions. However, there are many others, such as List and Pettit (2011), who think that groups can make decisions that cannot be reduced to decisions of sets of individuals. On this picture, sometimes the fact that *we* decided to φ can be an ineliminable part of each of our individual justifications for doing our bit towards our collective φ-ing. This too, along with the internalism about practical reasoning, is a viable, if contested, assumption in contemporary philosophy.

Against this backdrop, Collective Reasoning asserts that some friendships can be such that friends can rightly grant one another the power to contribute cooperatively with each's determination of their axiological commitments. That is to say, on this view we can allow our friends to share our authority to choose which values and ends to invest ourselves in. Accordingly, Blossius' declares that if Gracchus ordered him to do something – even to burn down the city's temples – then he would obey. According to Collective Reasoning, Blossius may be right to obey not because Gracchus' authority itself is higher than the importance of the civic and religious affairs, but because such an

order could only conceivably be one informed also by Blossius' own values. As such, the question of whether he would do what Gracchus might tell him to do amounts, for Blossius, to much the same as asking whether he would do what he himself might decide to do. For Gracchus' orders are thus not orders, but rather decisions made from the axiological perspective of Gracchus-and-Blossius qua friends and collective reasoners. The fact that it is they together that decide to φ can partly justify Blossius' decision to φ.

Collective Reasoning – and Montaigne's account of friendship which invokes it – have some *prima facie* attraction. Specifically, this view vindicates a sentiment about the importance that particular people can have for us as individuals. It allows that others can play a role in making us who we are that is not merely a causal role, nor merely an epistemic role. The appeal of Montaigne's view as I have interpreted it can be seen in connection with a well-known line from R.D. Laing's text *The Divided Self* (Laing 1960, p. 26):

> Here we have the paradox, the potentially tragic paradox, that our relatedness to others is an essential aspect of our being, as is our separateness, but any particular person is not a necessary part of our being.

Montaigne's view rebuts Laing's. For those people who truly have friends who have the status as collaborative reasoners in the formation of those people's basic axiological commitments, such friends qua 'particular persons' are, in a relevant sense, 'necessary for [their] being'. Collective Reasoning is attractive, therefore, because it averts Laing's 'tragic paradox'.

Despite its attraction, I do not think that Collective Reasoning can be true. The flaw lurking in this version of deep second-personality can be seen when one examines what might ground granting another person the status of co-author of one's own practical perspective. In a word, the problem that I would like to highlight is an echo of the necessary 'separateness' between persons that is also mentioned in the Laing quote above. This necessary separateness undermines the notion that a friend could ever truly be a co-author of the self, in the relevant sense. Ultimately, a self, or a self-identity, is the kind of thing that can have only one author.

Let us grant that one can identify with a friendship to such a degree that one values it more highly, and more unconditionally, than even one's most unquestioned ethical, political, or religious convictions. Such friendships are presumably uncommon, but let us suppose that they are possible. The problem for Collective Reasoning is that there is must be a distinction between the agential perspective of the subject who identifies so profoundly with their friendship, on the one hand, and the collective agential perspective that arises from that perspective on the other. That is, no matter how important a friendship is to an individual, there must be a self that is conceptually prior to the self that the friend collaboratively constitutes. And moreover, if any of the collective reasoning of the latter, socially moulded self is to be properly normatively grounded, then the value of the friendship must itself be normatively grounded to the prior, individual self. Thus, the reasoning that grounds the second-personal reasoning cannot itself be deeply second-personal.

To make this objection clearer, let's return to Blossius and Gracchus. Their friendship is so central to Blossius' self-identity that he would do whatever Gracchus would tell him to do. He would obey Gracchus not because he trusts Gracchus to know what's independently right, but because Gracchus' commands are spoken from the perspective of Gracchus-and-Blossius. Decisions made collectively by the two of them, that is, decisions made from their collectively constituted, shared perspective, have a weight for Blossius that none of his other convictions have. He identifies with their shared perspective to the extent that he might say that his own perspective just *is* their shared perspective, there is no difference between them. For this reason, decisions made by them both as friends have a bindingness for Blossius. The fact that the two of them have decided to φ itself counts as a reason for Blossius to do so, other things being equal, because Blossius grants sovereignty over his own deeds to his friendship with Gracchus.

However, this last sentence reveals the kernel of the flaw that undermines Collective Reasoning. It is *Blossius*, alone, who grants sovereignty over his own deeds to the joint-agent Blossius-and-Caius. And in order for the reasons that their friendship creates to be well-founded, Blossius himself must have good reasons for granting that sovereignty to their friendship. These could be reasons stemming from the value to his life of a friendship of that sort. Whatever they are, though, they cannot themselves be deeply second-personal, on pain of regress. And for that reason, neither are the reasons built out of their co-constituting friendship really deeply second-personal. They are not second-personal all the way down, since they are ultimately grounded in these non-second-personal reasons (such as the value to Blossius of the friendship).

Conclusion

At least since Darwall's work leading up to *SPS*, the concept of the second-person relation, and of second-personal reasons, have been discussed a great deal in practical philosophy. There has, however, been a regrettable degree of ambiguity in these discussions, where a number of different senses of second-personal reason have been elided. In this chapter, I have offered a schematic way to delineate these senses in the hope of making it possible to clarify much of that interesting and important philosophical conversation about the nature of the practical influence of second-persons as such.

In addition to merely clarifying things, I have advanced the view (to my knowledge a novel one) that the second-personality of a reason is something that can come in degrees. I have distinguished between different degrees of second-personal reason and offered a set of arguments animating the plausibility of many of these kinds of reasons being real features of the normative landscapes within which we live our lives. However, in these last sections, I have also argued that it is a mistake to think that some reasons are second-personal all the way down, or at least that much more would need to be shown to make such a view a live option.

Chapter 4

The discretionary normativity of requests[44]

In Chapter Two I argued that moral demands are not the only the kind of reason which depend for their normativity on second-person relations. (Indeed, I further argued that moral demands may not even be second-personal reasons themselves.) I have since sketched the conceptual terrain for a number of other kinds of practical reason which are all, on the face of it, second-personal. Of these, requests stand out as a particularly interesting site in our practical lives at which it might be possible, through analysis, to gain a better understanding of how it is that our mutually recognitive relations with one another can give rise to normative forces that govern what we ought to do. The particular intrigue of requests resides in the way that they draw somehow on the relationships between requester and requestee to create reasons which nonetheless preserve the discretion of the requestee to decline doing what they have been asked to do. If my arguments have been successful in showing that neither moral demands, nor deep second-personality of the self, nor deep second-personality of the other, are really sets of reasons that are second-personal all the way down, then requests may be exemplary of the kind of reason that are, more than any other, determined by second person relations.

Why should we ever accede to requests? Is there some common normative basis for the reasons that we act on when we do accede to requests, and if so, what is it? One motivation for trying to answer these questions is the desire to vindicate requesting as a *rational* form of interpersonal reason-giving. This is a necessary piece of the broader vindication of our deeds as those of rational agents. If there is no good answer that explains the normative foundations of requesting, then there is a sort of nihilism hovering in wait. That is, if there are no sound normative grounds for acceding to requests, then perhaps we only ever do so out of socially inculcated habits which themselves cannot be justified. Interestingly, scholars in the field of cross-cultural pragmatics have found evidence to support the view that our tendencies to use requesting as a form of reason-giving are culturally-relative phenomena.[45] Perhaps this view itself supports the aforementioned nihilism about requesting: we only make and accede to requests out of culturally contingent habits, not on the basis of normatively grounded reasons. So there is philosophical worth in seeing what defence can be made against that nihilist perspective.

---

[45] See for instance (Blum-Kulka et al. 1989), (Economidou-Kogetsidis 2010), (Huangfu 2012).

A second, further motivation for thinking about the normative basis of requests comes from its possibly fruitful connection to what some philosophers call the second-person relation.[46] If, in certain circumstances, we can create, at will, genuine reasons for others to act, then what does this say about the kind of relations that we always stand in to other people? One possible thing it might say is that just by virtue of being a person, one has the status in the eyes of others as a source of practical reasons. More than that: as one whose intentions for other people to act can themselves count as good reasons for those others to so act. If it is indeed *just* by virtue of being a person that one has the power to make a request of another, then it might look as though the normativity of requesting is a structural feature of the second-person relation. That, I take it, would be an interesting conclusion, and would also have interesting ramifications for moral philosophy. It is not exactly the conclusion that I will defend here, though the intrigue of this hypothesis is part of my motivation.

The thesis that I *am* going to defend is this: a request requires for its efficacy that the person addressed by the request (the addressee) places discretionary value in the person making the request (the addressor). My primary task is to present and stand up for that thought, so for the purpose of this chapter it will only be between the line that I can make conjectures about the consequences of this view for thinking about the second-person relation (a matter that I will revisit explicitly in the Conclusion to the dissertation). This thesis is driven by a need to cover the following three explananda, which I believe set the bar for any attempted account of the normative quality of requests. (i) Requests can in principle be the source of *legitimate* reasons; (ii) requests create *new* reasons; and (iii) the reasons that they create are in some sense *discretionary*. To be sure, this triad is, at least at first glance, hard to reconcile.

However, that is the task at hand and it will be undertaken in the following manner. First I will discuss the definition of requesting, making clear how it is distinct from other kinds of interpersonal reason-giving. In the second section, I will set out the notion of a discretionary value, explaining the particular sense in which one could place discretionary value in a person. Third, I will lay out a theory by David Enoch (Enoch 2011) of what he calls 'robust reason-giving', which is a class of reason-giving that includes requesting. I am sympathetic to Enoch's picture and I will try to illustrate its advantages, but I will also argue that it fails to account for the whole triad of explananda when it comes to requests. As such, in the fourth section, I will propose a solution – an adaptation of Enoch's theory to specifically explain the normative structure of requests. There is an apparently compelling objection to my proposal that I will address in the fifth section. If my thesis holds true then it provides a rebuttal of the nihilist suggestion that requests can never create well-founded reasons, and a vindication of requesting as a feature of interpersonal life.

---

[46] I am thinking here in particular of Stephen Darwall (2006), though the interest in the second-person is now much broader, as illustrated by two recent special journal issues (Conant & Rödl 2014b), (Eilan 2014).

1. Distinguishing requests

I shall define requests as follows:

> **Definition:** A request is an attempt by an addressor to create and communicate a non-obligatory reason for the addressee(s) to perform an action.

A successful request is thus one that succeeds in this attempt. Since my goal in this chapter is to consider the reasons with which we are presented in requests and to identify their normative basis, defining what exactly requests *are* is a separate matter, though a crucial one. As such, for the purpose of my main argument, this definition is stipulative. Having said that, I do hope that the concept of requesting under inspection here rings true as a familiar device in the normative play of interpersonal relations. There are, perhaps, a few points in my definition that could be contested, so in this section I will briefly defend two of the most salient of those points. That is, I will defend the ideas that requests should be conceived as *creating* reasons, and that those reasons are *non-obligatory*.

One sceptical perspective from which one might criticise this definition is that of ordinary language. One might reasonably object that when we talk about requests, we are often talking about utterances that do not create reasons but merely state reasons that were already in play; or similarly, it may be that we use the term 'request' to refer to the exchanging of reasons that are obligatory, not non-obligatory. I do not contest that as we commonly use the term, it does often include these features that are not captured by my definition. I am not providing a definition of the concept as it is used in ordinary language. Rather, the goal here is to define requests as a form of reason-giving with *distinctive normative force*. Specifically, that distinctiveness from other forms of reason-giving derives from thinking of requests as uniquely occupying a quadrant in the chart below (fig.3). Again,

Create new reasons

Commands	Requests

Obligatory   Non-Obligatory

Demands	Epistemic reasons (telling, advising)

State prior reasons

73
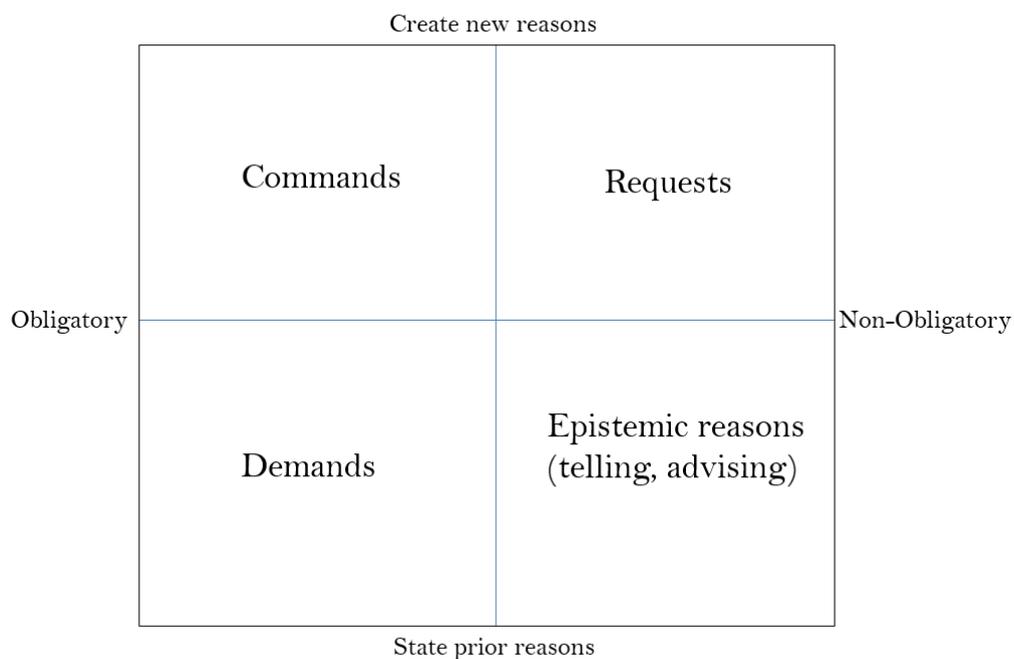
whilst this definition is stipulative, I also believe that it captures the heart of the concept of requesting. Thus, when in ordinary language we call something a request though it does not meet these conditions, I suggest that we are thereby deviating from, and ever so slightly perverting, the true meaning of the term.

Figure 1 illustrates the conceptual terrain for interpersonal reasons – the kind of reasons that one person is able to give to another person, intentionally, in an act of address, thereby applying some rational force for or against any particular action.[47] The chart carves the terrain along two dimensions: whether the reason presented is thereby created, and whether it is obligatory. As such, the top-left quadrant represents interpersonally-given reasons that are both created and obligatory, which I suggest might be the conceptual heart of the notion of a command.[48] Beneath that, in the bottom-left, are reasons that are also obligatory, but which are not newly created. In these instances, the addressor presses their addressee with an obligation that putatively befalls the addressee irrespective of this instance of it being addressed. Again, my suggestion is that this is the conceptual heart of the notion of demand.[49] In the bottom-right are reasons that are also not newly created, but which are not obligatory either. This is just a kind of purely epistemic reason-giving, as when one tells another some non-normative fact ('the bus is coming in 5 minutes') or a normative fact ('the item you have in your hands is very precious'). These are *pro tanto* reasons, considerations which exist, and pertain to the addressee, irrespective of this particular interpersonal address. Finally, then, in the top-right quadrant are non-obligatory reasons that are newly created in the act of being addressed to one person by another: requests.[50, 51, 52]

---

[47] Note that in categorising ways of giving reasons for action, this chart does not purport to categorise ways of giving reasons for belief.

[48] More precisely, I mean two things: that the conceptual heart of 'commanding' is the act of intentionally creating an obligatory reason for another person; and also, broadly, that commanding is a paradigm representative of the class of acts which create obligations.

[49] One might be suspicious of this demarcation in Figure 1, on the basis that the border between demands and commands is rather blurrier than I am suggesting: that sometimes demands create new reasons. Let me try to allay this concern through an example. Consider a group of employees in a factory that produces supermarket sandwiches, who, after a breakdown in official negotiations, down their aprons and march to the management corridor of their employer's head office, where a spokesperson for the workers issues a demand to the employers that they introduce a decent workplace pension scheme. One way to interpret this putative demand would be to understand it as stating an obligation that was already in play, and thereby holding the addressee to account, much like the rebuke one might give retrospectively by saying 'you should have provided a pension scheme for your workers'. Indeed, the spokesperson's utterance of the demand could at the same time be an instance of epistemic reason-giving, making the employers aware of this normative fact that they had hitherto overlooked. Alternatively, one could understand this 'demand' as being, in fact, a command in disguise. That is, the spokesperson may be creating an obligation simply by stating it, invoking their authority as representative of the people on whom the employers depend. Either way, it is apparent that the distinction between demands and commands is a clean, not a blurry one: either the obligation that the utterance purports to present was already there, or the utterance purports to create it. Little hangs on the terminological dispute over what we use the words 'demand' and 'command' to refer to. The point for my purposes is that these terms can be used to capture two *discrete* normative moves.

[50] One of the few philosophy papers that expressly sets out to address the normativity of requests provocatively argues that requests *do* create obligations. That is Cupit (1994), who thinks that requests appeal to obligations that are grounded in the commitments of the agent. Despite appearances, this view is interestingly similar to the one that I will develop here, though I do ultimately disagree about the characterisation of the resulting reasons as

An example here may be helpful in illustrating the extent to which this stipulative definition is revisionary of the ordinary concept of requesting, but also the extent to which it nonetheless succeeds in capturing the core of that concept. Consider Bronwen, a bus mechanic who has recently taken on young Mair as an apprentice. Bronwen is kindly and warm-hearted towards Mair and I suspect that most people would ordinarily take her to be making a request when she asks Mair, 'Please would you change the oils of the Leyland National that came in yesterday?' But if Mair were to treat this reason that she has just been given as *discretionary* (a notion that will be elaborated upon shortly), then I fear that Mair would have misunderstood the situation quite badly. If she simply chooses not to change the oils of the Leyland National then she will have something to answer for, and Bronwen could legitimately express at least a little irritation towards her.

Similarly, suppose that a few hours later, Bronwen notices that Mair has not gone near the Leyland National, and so she reiterates her earlier utterance. Again, there is an ordinary sense in which the reiterated utterance would commonly be thought of as a request. But woe betide Mair were she to think of utterances like these as creating *new* reasons that were not in play prior to their being uttered. Whilst we may ordinarily refer to both Bronwen's first and second utterances as requests, on my stipulative schema, they are a command and a demand, respectively. The first created a reason for Mair to change the oils that she didn't have prior to having been asked, but a reason that is obligatory, rather than discretionary. The second did not create a reason at all, but merely reminded Mair of an obligation that she had been given, where the reminder itself was an act of holding Mair to account to that obligation.

Whilst being revisionary, though, there is nonetheless some purity to the notion of requesting that is advanced here. The claim that the given stipulative definition captures the conceptual heart of requesting rests on the further intuition that although we may ordinarily refer to Bronwen's speech acts as requests, we may at the same time acknowledge a sense in which they were not *really* requests. I make this sentiment explicit by claiming that speech acts are only *really* requests when they attempt to create non-obligatory reasons, and stipulate that hereafter in the present discussion, the term 'request' will be reserved for the real deal, so to speak. Thus, when appeal is made below to intuitions about requests, these appeals seek to draw on intuitions about only this central subset of what the term 'requests' often includes.

I have thus put aside the ordinary language objection on the basis that what is required here is a stipulative definition that captures the distinctive normative force of requesting. But that will not

obligations, rather than discretionary reasons. Also noteworthy is that Cupit (1994, p. 449) agrees with the other half of my definition: that requests 'generate' new reasons.

[51] Defining requests as giving non-obligatory reasons is concordant with Lance and Kukla (2013, p. 460): 'The [normative] output of a successful imperative is an obligation on the part of the person ordered to do what the speaker ordered her to do. The output of a successful request is that the target now has a specific sort of reason to do what was requested, but it is essential to the notion of a request that this reason is *not* an obligation.'

[52] This definition of requests also conforms to Raz's understanding of requests as a kind of content-independent reason-giving (Raz 1988, pp. 36-7), with which Owens (Owens 2012, p. 86) also concurs.

satisfy a detractor who maintains that not only does the word 'request' commonly refer to reasons in the other three quadrants of the chart here, but moreover, that the quadrant that I am designating to requests is either unnecessary or an impossibility: there can be no such things, and/or we need not think that there are such things, as speech acts that create non-obligatory reasons. In a way, the thesis that I want to defend in the later sections of this chapter speaks to this challenge by offering an explanation of how there *could* be a normative basis for such reasons. But a more immediate response to this challenge can also be made.

The challenge in question, to be clear, rejects my stipulative definition of requests on the following basis: that the other three quadrants of the chart between them exhaustively explain the reason-giving that goes on in requesting; that no recourse is needed to the idea of a newly-created non-obligatory reason. I think this is not true. An example might help to illustrate the fact that the kinds of reason-giving categorised in those other three boxes are inadequate to explain at least some of what we conceive ourselves to be doing when we make and accede to requests.

Consider two friends, Sioned and Ffion. Sioned is mounting an election campaign and she wants Ffion to help as her campaign manager. Committing to the campaign would constitute a substantial sacrifice for Ffion as it will be stressful, and for the course of the campaign it will take a lot of time away from her own work, her family and her other engagements. Suppose that Ffion knows perfectly well that Sioned wants her help: indeed, everybody knows it. But because of the extent of the sacrifice that it would entail, Ffion has not voluntarily offered her help to her friend. For some time, Ffion knows that Sioned desires her help and Sioned knows that Ffion knows this too, but somehow she cannot bring herself to ask for help: partly out of pride, partly out of reluctance to burden her friend, partly in the hope that an offer will be forthcoming from Ffion anyway. But it is not, so the time comes and Sioned confronts the awkwardness that has arisen between them with a request: she explicitly asks Ffion whether she would commit to helping Sioned's election bid in the role of her campaign manager. This, I suggest, is a request which presents a non-obligatory reason for action. But moreover, the request itself has altered the normative situation. I suggest that it has done so by creating a reason that was not present before.

All I mean to appeal to here are two intuitions about this case (and therefore about others like it). The first intuition is that the reason presented by the request is not an obligatory one. Of course friendships are relationships that can sometimes generate obligations, including obligations to help one another.[53] But being someone's friend does not mean *always* being obliged to help them, regardless of what the help is needed for, or of the cost that helping would incur. Sioned's request is supposed to be an example of one of those deeds which Ffion is not obliged to do: Sioned is not at risk of being harmed and the costs of helping are substantial. As such, the kind of reason-giving that is

---

[53] For a discussion of how friendships and other particular relationships can generate special obligations, see (Jeske 2008).

going on in this instance cannot be anything on the left-hand side of the chart in figure 1, since the reason is non-obligatory.

On the other hand, the second intuition is that the making of the request is a normatively significant moment. Whatever the act of uttering the request does, it has some kind of impact on the balance of reasons for and against helping. That is to say, from Ffion's perspective it makes a difference that Sioned asks for her help. Since in the example it is stipulated that prior to the request being made, Ffion already knows full well of Sioned's desire for her help, it seems that the normative difference that the speech act of the request makes cannot be an epistemic matter. It cannot be the case that Sioned alters the normative situation through some epistemic reason-giving, telling Ffion about her desire for the latter's help, because Ffion already knows all about it.[54] As such, the normative role that the request is playing must be more than merely epistemic. Somehow or other, this utterance is creating a new reason for Ffion to commit to the campaign itself.[55] When added to those of which she was already aware, this utterance could conceivably be sufficient to tip the balance of reasons for Ffion in favour of making the commitment to her friend. This prompts the question that I will try to answer later on, of *how*, without any appeal to authority, it is possible for people to simply create reasons of this sort for others.

Despite having said earlier that the definition of requesting in this chapter is not intended as a descriptive attempt to capture the ordinary language meaning of the term, I have nonetheless addressed the relation of my stipulative definition to the ordinary language conception. Whilst in everyday speech we may commonly use the term to refer to other forms of reason-giving, I have argued – through the example of Sioned and Ffion – that the everyday notion must at least *include* the kind of reason-giving that is under inspection here. In the next section, I will elaborate on the idea of having reasons that are discretionary, and how other people could be the source of such discretionary reasons. This will lay the ground for my actual account of the normative basis of requesting, which I will present later.

2. The discretionary value of persons

What it is to value something is a troublesome question. But whatever else it is, valuing X is having a favourable attitude toward X such that X can be the source of reasons for certain actions. Thus, if you value this photograph of your grandmother, as well having an array of beliefs about, and emotional attitudes toward it, you will also be disposed to act in certain ways regarding the photograph. You

---

[54] As I will discuss in section 5 below, some theorists – particularly in the domain of speech act theory – do think of requests as functioning by expressing desires. (Searle 1969) is the precedent for this.

[55] One objection here to the claim that requests create new reasons, would be to point to cases where the addressee seemed to have more than sufficient reason to perform the action in question prior to being asked, where the addressor could legitimately say 'I should not really have to ask!' However, in such cases as this, it is at least *possible* that the request does indeed add a new reason to the existing pile. It is exactly the fact that the existing pile *without* a further reason had not already moved the addressee to action that has drawn the addressor's ire.

may, for instance, be prone to lurch to catch it if you were to see it falling; or perhaps you would be inclined to act to ensure that it retains a prominent, visible position on the mantle-piece – whatever.[56] In this section I want to set out the sense in which valuing another person can make them the source of reasons in this sense (that is, in the internalist sense that they would not be such a source of reasons were it not for one's attitudes towards them). It will also be helpful here to clarify what bearing this kind of internalism about this class of reasons has on other important matters, namely, the non/obligatory quality of the reasons, their relation to external reasons, and to moral reasons.

A first important thing to note is that we are morally obliged to value others, to a certain extent. Typically, we might express this by saying that we are obliged to respect others, or some similar thought. It seems plausible to me (though nothing hangs on this here), that our obligation to respect all people does not simply pertain to performing the deeds of respecting their dignity. Rather, the obligation also pertains to the attitudes that we hold. We must actually *respect* others' dignity, which means, *inter alia*, believing that their dignity is worthy of respect and perhaps harbouring at least some minimal degree of emotional connection with their dignity – a disposition to regard threats to their dignity with concern, for instance. Holding the attitude of respect toward others is holding an obligatory interpersonal valuing-attitude.[57]

Indeed, respect might be an attitude that we are each obliged to hold toward everyone else, but there are other obligatory interpersonal valuing-attitudes, ones which are specific to one's particular relationships with others. Often, for instance, one is obliged to hold the attitudes of deference towards one's elders, of sympathy towards one's young children, of solidarity with one's friends and comrades. In entering into relationships with others, one accrues obligations to hold certain attitudes toward them. And the obligatoriness of such attitudes is not undermined by the fact that forming those relationships with those people was not, in the first place, something one was obliged to do.

But not all interpersonal valuing-attitudes are obligatory. Consider the affection you might have for your nieces, the admiration you might have for a colleague, pity for an unfortunate stranger, endearment to a charming one, the lofty esteem that you might have for an able rival. These attitudes are all discretionary. No-one is entitled to demand them, or to blame others for preferring not to hold them. In fact, for any valuing-attitude that can be legitimately expected, one can imagine the possibility of holding that attitude more intensely than is expected. The fact that one is obliged to hold a valuing attitude to a certain extent implies that whether to hold it to a greater extent is a matter of discretion.

---

[56] See (Scheffler 2012, pp. 27-8).
[57] I do not think that one is obliged to *always* hold this attitude toward *everyone* in the world. Whilst there may well be obligations that we all owe to absolutely everyone, the obligations to hold certain attitudes toward people only emerge when you are (or should be) aware of those people.

Significantly, there is a difference between valuing a quality, or set of qualities that someone has, on the one hand, and on the other, valuing *them*.[58] When discussing interpersonal valuing in requests – which I will get on to in the next sections below – the kind of valuing must necessarily be the latter kind, not the former. Various qualities may play a role in leading us to place value in others, or in vindicating in our own eyes the value that we already do place in others. But our interpersonal relationships are characterised by the way in which we value persons themselves, not merely certain aspects of them.

A consequence of the fact that it is people as individuals that are the objects of our interpersonal valuing-attitudes is that the reasons that arise from those attitudes are reasons to act *for others*. Thus, when my admiration of you manifests in giving me a reason for some kind of action – to help you in your endeavours, perhaps – I thereby have a reason to act *for you*. Doing someone a favour – which is a way of acting on an interpersonal valuing-attitude – entails doing something *for* they themselves. What this kind of inter*personal* valuing attitude contrasts with would be a kind of valuing attitude that pertains only to particular traits and qualities: fondness for Harriet's dry wit but no fondness at all for Harriet.[59] The point here is one that will be relevant later: that there is a species of discretionary valuing-attitudes that play a central role in our social lives, which have as their objects people, as such, rather than merely the valuable qualities that people sometimes bear.

At this point, it is worth noting the distinction between being obliged to do something, on the one hand, and having *reason* to do something all things considered.[60] That is, it might be said that one 'rationally ought' to do that which one has most reason to do, all things considered. But in distinguishing a set of attitudes that one must hold, from another which one may hold at one's discretion, I am not making any claim about what one rationally ought to do, or where rationality allows for some discretion.[61] Rather, the sense of obligation at play here is to do with what it is morally right and wrong for us to do.[62]

---

[58] The difference I want to capture matches Darwall's (1977) distinction between appraisal respect, and recognition respect, respectively. Having said that, I do not mean to endorse what Darwall thinks are the bounds of the recognition respect that persons can enjoy from one another, where 'there can be no degrees of recognition respect for persons' (1977, p. 46). Unlike Darwall in that paper, I am concerned here with *non*-moral interpersonal valuing, within which sphere it is possible to value some people more than others, in a way that is not at all a matter of appraising their virtues.

[59] The valuing I have in mind could be characterised as valuing someone *de re* and not *de dicto*. For a discussion of related matters see (Kraut 1987, esp. p. 423).

[60] Thanks to Bob Stern for pressing this distinction.

[61] As it happens, I do also in fact think that there is such a thing as rational discretion *and* that interpersonal valuing-attitudes that are discretionary in the deontic sense are also discretionary in the sense of there being no determinate all-things-considered set of attitudes that any given person rationally-must adopt. But for present purposes, the notion of rational obligation is not relevant.

[62] What characterises obligations is a matter of some controversy. One influential account is Raz's view of obligations – or mandatory reasons – as involving a second-order 'exclusionary' component that instructs the disqualification of competing first-order considerations (Raz 1999, pp. 73-76). Other theories of obligation define the concept in terms of the kind of accountability that it implies (Darwall, 2006, chpt.5). For my purposes, I need not endorse one account or another, so long as they are all compatible with a general thought that

Moreover, the obligatoriness of a valuing-attitude finds expression in the obligatoriness of the actions that express that valuing-attitude. I have already mentioned the strong connection between valuing something and treating it as a source of practical reasons. This connection illustrates the divide between obligatory and discretionary interpersonal valuing-attitudes. Suppose that Charlene is obliged to hold an attitude of deference toward her professor, Dominique. If she fails to hold that attitude she will be doing something wrong. This failure may manifest itself in certain deeds – an insufficiently deep bow, a lacklustre display of courtesy, perhaps – and by extension, these deeds too are wrong. By contrast, there is nothing wrong about failing to feel heart-felt affection for someone, or genuine admiration, or real pity. It is intrinsic to the very notions of these attitudes that one cannot be obliged to feel them. The true sentiment of affection can only be an organic sentiment, one that arises naturally and not out of duty.[63] The same goes for the 'true sentiments' of other interpersonal valuing-attitudes.

The notion of a valuing-attitude being discretionary is crucially important. An attitude is discretionary just when it is not obligatory. And on the picture of practical reasoning that I am assuming in this chapter, an attitude is non-obligatory just when no-one is entitled to react with anger to one's holding or failing to hold the attitude. When we are unconstrained by duties, in this sense, we must exercise discretion over our conduct: we take ownership over which values to invest ourselves in, and over how we weigh those values against one another. This thought too will be relevant later in the discussion of requests. Since requests do not create obligations, I will claim that they must make an appeal to their addressees as agents who have this kind of discretion.

In this section, I have been trying to express an idea that I think comes naturally when we think about the reasons that we have to act in the interests of others. That is, I have tried to establish – in line with common intuitions – that there are such things as discretionary reasons to act for another person that stem from discretionary attitudes of valuing that person. Before moving on, there are two noteworthy features of the general picture of practical reasoning to emphasise. The first is that, as mentioned above, the reasons that I have been chiefly concerned with are 'internal' reasons. That is to say, they are normative reasons that make essential reference to some aspect of the motivational set of the agent for whom they are reasons.[64] However, that does not mean that the theory that I am advancing represents a partisan position on the debate in meta-ethics between internalists and externalists about moral reasons; it doesn't. It is entirely compatible with the view that people have some reasons that depend on their own values, to also think that they may also have some other reasons – moral obligations, perhaps – that are external to their own set of values. I thus remain

---

obligations 'always give agents conclusive reasons for acting that outweigh or take priority over any potentially competing considerations' (Darwall, 2006, p. 26).

[63] For a further defence of the view that there cannot be a duty to love, see (Driver 2014). Not everyone holds this view, however, see (Liao 2007).

[64] For an elaboration on this kind of understanding of internalism, see (Markovits 2014); for its classic source, see (Williams 1981, p. 102).

neutral on that question. Having said that – and this is the second noteworthy point – the kinds of discretionary other-regarding reasons that I have identified *could* have some moral significance. Specifically, when moral philosophers talk of something being the wrong or the right *kind of reason* for someone to do something, the discretionary interpersonal reasons discussed here might seem relevant. It seems plausible, *prima facie*, that acting out of a genuine, discretionary heartfelt desire is very much the *right* kind of reason to act, even if the action itself is something that one is morally obliged to perform. So this is just to note that whilst the picture that I am advancing is neutral between competing moral and meta-ethical theories, it may have some interesting consequences.[65]

3. Enoch's account

It is now possible to return to the goal of explaining what, if any, normative force there might be to the new, non-obligatory reasons that are presented in requests. David Enoch (2011) has devised a sophisticated account addressing this issue. Here I will offer a sketch of how Enoch proposes to explain the normative power that people have to give practical reasons 'robustly'. There is, though, a problem with this explanation when it is brought to bear on requests. I will try to illuminate the difficulty that Enoch's view has in accounting for the discretionary quality of the normativity of requests. In the following section below (section 4), I will propose an amendment to the view that enables it to overcome this problem.

The primary concern driving Enoch's inquiry is the explanation of a phenomenon broader than merely requesting. The motivating question is rather, 'if, as seems likely, 'reason must constrain and guide the will', how is it that we can create reasons at will'? (Enoch, 2011, p. 1) The sphere of intentionally created reasons includes commands and promises, and to the entire domain he gives the label 'robust reason-giving'. But Enoch thinks that requests are the paradigm form of robust reason-giving by virtue of their simplicity relative to these other apparently more complex phenomena, where authority is involved. As will become apparent below, requests have complexities of their own. But the theory of robust reason-giving in its general form is still of use. According to that view, all practical reasons that one person can give to another can be categorised into two sets. On the one hand, they could be a kind of merely epistemic reason-giving (the sort of thing found in the bottom half of the chart in figure 3 above). Alternatively, if they are doing something more than merely telling or advising the addressee about some prior existing reasons, then, Enoch thinks, they must be *triggering* a reason. As such, since these reasons by stipulation are more than merely epistemic reasons, they must be instances of triggering-reasons.

It may be helpful to elaborate on this point. A robustly-given reason – such as a command or a request – does its normative work not simply by trying to reveal to the addressee what reasons there

---

[65] Interesting though such consequences may be, discussing them properly must remain a matter for another time.

are for them to act, but by in some way *changing* what such reasons are. But practical reasons are not the sorts of things that can be merely willed into existence wantonly. You cannot, for instance, make it the case that a stranger should arduously undertake to do your bidding, merely by deciding that they should. Rather, Enoch (2011, p. 9) infers, these reasons work by *realising the non-normative antecedents* of conditional reasons that hold true independently.

One example of this can be seen in the case of commands, which are a species of robust reason-giving. When the sergeant commands one of her officers to quick march to the barracks, she triggers a conditional reason, by realising its antecedent. That conditional reason must have the form: 'if commanded to do so by the sergeant, then the officer has (obligatory) reason to quick march to the barracks.' And according to Enoch the same story applies to requests. When Sioned requested Ffion to help her with the campaign, she triggered something like the following conditional reason: 'if requested to do so by Sioned, Ffion has (non-obligatory) reason to help with the campaign.' Such reasons can be made when the relevant conditional reasons are true; conversely, successful robust reason-giving implies the truth of the prior conditionals (Enoch, 2011, p. 10). [66]

This is not the whole of the account, however. So far, no space has been made for the difference between robust reason-giving, on the one hand, and on the other, the variety of other ways in which non-normative circumstances can be manipulated so as to trigger conditional reasons. Enoch (2011, p. 4) gives the example of the neighbourhood grocer raising the price of milk. By doing so, we can suppose that she triggers (again, by realising the antecedent of) a prior conditional reason, one of the following sort: 'if the price of milk at this shop is above X amount, it is too expensive so you should not buy milk here.' Here, the mechanism by which the grocer *inadvertently* gives you a reason not to buy her milk looks identical to the mechanism by which the sergeant creates a reason with her command, or that by which Sioned creates her request. However, this is inadequate since it certainly seems that the normative power of reasons like those presented in requests simply have more to them, so to speak, than such incidental reason-giving as the grocer's price change creates.

To address this, the account of robust reason-giving must incorporate the role played by the *intentions* of the parties to these exchanges. In requests and commands, the reason that one attempts to make with one's utterance depends on the addressee recognising one's intention for this utterance to give them a reason. Enoch's exact formulation of this thought summarises the account (Enoch, 2011, p. 15):

> One person A attempts to robustly give another person B a reason to φ just in case (and because):
> (i) A intends to give B reason to φ, and A communicates this intention to B;

---

[66] Following a line of argument by Mark Schroeder (2006), Enoch (2011, p. 11) acknowledges the conceptual possibility of robust reason-giving that does not trigger a prior reason but really creates a wholly new reason. Thus, divine command theorists may believe that the obligatoriness of a command *consists* in God having commanded it. On such a view, God gives reasons that do not rely for their force on the truth of prior conditional reasons. But in the present chapter I am concerned only with reasons exchanged between ordinary mortals, which, contrarily, *must* be grounded in prior, conditional reasons.

<div style="margin-left: 2em;">

(ii) A intends B to recognize this intention;

(iii) A intends B's given reason to φ to depend in an appropriate way on B's recognition of A's communicated intention to give B a reason to φ.

</div>

There are several considerations to be discussed in relation to this proposal that I shall leave aside here. For the purposes of my argument, I shall assume that the formulation of robust reason-giving set out here sufficiently explains the general manner in which intentions are relevant to the class of normative interpersonal interactions that are at issue – including commands and requests. In what follows, I will focus in on requests and inspect the status and normative quality of the conditional reasons – those that are required for a request to be successful, whose antecedents are made-true by the uttering of a request.

At this juncture, I would like to raise a problem. Or rather, to raise again the problem with which I began. How does the theory of robust reason-giving fare at dealing with the three explananda of requesting? The first explanandum was that requests can in principle be the source of *legitimate* reasons. One worry in this regard might be that in acceding to requests we act out of socially inculcated habits, rather than on the basis of well-grounded reasons. Another concern is that requesting must always be some kind of coercion, since legitimate reasons cannot simply be willed into existence, out of thin air. The theory of robust reason-giving addresses these concerns. It does so by revealing the role of prior conditional reasons that are brought into play in requests. Those are reasons of the form, 'person A should φ, if requested to do so by person B'. These conditional reasons are not willed into existence, they are in some sense there already, before – or at least *at* – the moment of a request being made. By positing the existence of such conditional reasons, the theory can explain in principle how acceding to a request could be justified by reference to a legitimate reason. Acceding therefore need not be thought to be a response to a social convention or being bent coercively by the mere will of the requester.

The second explanandum was that requests create *new* reasons. This is the feature towards which the theory of robust reason-giving is primarily addressed. Enoch is motivated by a suspicion of the mysterious sounding notion of reasons – which bind our wills – coming into existence at the whims of mere agents. But this mysteriousness is played off against the phenomenology of requests. As the example of Sioned's request illustrated earlier, it seems certain that it is possible for requests to make an impact on the normative terrain – to *do* something, that is – even when all the relevant normative and non-normative facts are known, so they cannot be *doing* anything epistemic. Robust reason-giving explains this *doing* as kind of a triggering. By making-true the antecedent of a prior, conditional reason, requests manipulate the non-normative circumstances in such a way that the addressee has a reason that they didn't have before. (Namely, that reason is just this: that they have been requested.)

The third of the triad of explananda was that the reasons that requests create are in some sense *discretionary*. How does the theory of robust reason-giving account for this discretionary quality? This is where I think the problem arises: I do not think it can. For a request can be made only if there

is a prior conditional reason available for it to trigger. In other words, I can only request you to φ if it is the case that 'you have reason to φ, if I request you to do so'. If such a prior conditional is not true, then the request will fail. It will fail not just to persuade the addressee conclusively to φ, but even to alter the balance of reasons at all. But on the other hand, if the prior, conditional reason is true, and it is triggered, then the addressee simply has a reason to φ, and it is not clear where the discretionary quality enters in.

This is quite a serious concern. As we have seen, it is a fundamental quality of requests that acceding to them is distinctively a matter of discretion. Asking someone to do something is an interesting, special form of reason-giving precisely because in so asking, one intends for the other to treat this request as a reason, but not for them to treat it as itself conclusively instructing them. We ask them to φ, and thereby acknowledge that whilst our wishes are clear, the matter of whether to φ or not is up to them. Despite the fact that Enoch thinks of requests as the paradigm of robust reason-giving, that theory lacks any conceptual resources to account for this defining discretionary quality.

It might be thought that Enoch's account as it stands can accommodate the discretionary quality of requests simply by distinguishing them from commands.[67] That is, the reasons presented by requests, rather than being obligatory, are merely *pro tanto*: they are 'first-order' considerations that favour certain actions, but they could just as well be outweighed by more pressing reasons that speak against those actions. Obligations should not be outweighed in this way. Obligations purport to provide conclusive reasons. Thus, the discretionary quality of requests might be thought to consist simply in the fact that they do not purport to provide conclusive reasons.

But this line of defence, though initially tempting, fails on two fronts. First, it simply begs the question. It is true that to be merely *pro tanto*, rather than obligatory, is part of what it is for a request to be discretionary. But merely claiming that requests trigger conditional reasons and those reasons are merely *pro tanto* is not enough. The question is *how* can one person issue another with a merely *pro tanto* reason to do as they ask? In the case of commands, the prior, conditional reasons ('if the sergeant commands, then the private should φ') are made-true by the authority that the addressor has and the reasons are grounded in the normative grounding of that authority. These normative powers are often taken for granted: the parent's authority over their teenaged child plausibly stems from their parental responsibility; the employer's authority over their employee stems from an explicit contract; the restaurant customer's from an implicit one; and so on. To be sure, in any given case, it may be far from straightforward to determine whether the putative authority is in fact well-grounded. What is straightforward, though, is that *when* reasons are robustly-given in commands, they are *always* grounded in this kind of authority. In requests, though, there is no authority at play. So what, on Enoch's picture, could ever make it true that person A has a merely *pro tanto* reason to φ if requested to do so by person B? This is the question that the theory is so far ill-equipped to answer.

---

[67] Thanks to Daniel Viehoff for pressing this point.

The second front on which that tempting line of defence fails is its characterisation of the discretionary quality of requests. One of the distinctive things about requests is that, at some level, they appeal to addressees to *choose* what to do: to make a choice between competing values, rather than simply to calculate what they have most reason to do. Or, put another way, when we do appeal to others to make a certain choice between competing values, it is possible for us to make such appeals by requesting. But insofar as requesting is understood merely as the simple triggering of a reason, this aspect of the phenomenon remains mysterious. The general theory that Enoch has articulated has no resources to explain how requests can appeal to their addressees to choose between such competing paths where only they, the addressees, have the authority – the discretion – to make that choice.

4. A proposal

In response to these difficulties, I have a sympathetic proposal to amend the account of robust reason-giving as it applies to requesting. My suggestion is that the prior reasons that are triggered by requests must always have not one but two conditions. Besides being conditional on a request being made, they must also depend on the addressee placing some discretionary value in the addressor. Accordingly, those prior, conditional reasons take the following general form:

> Person A has a reason to φ *if* [condition (i)] requested to do so by Person B *and if* [condition (ii)] Person A places sufficient discretionary value in Person B.

On this view, then, the normativity of requests is keyed to interpersonal valuing-attitudes. Specifically, it is keyed to a set of valuing-attitudes which are a matter of an agent's discretion: these are attitudes that one is not obliged to hold. Specifying this point, therefore, introduces sufficient conceptual resources to explain the discretionary quality of the reasons presented in requests.[68]

My proposal bakes-in the discretionary character of the reason at the level of the prior, conditional reason. A consequence of doing so is that if that discretionary valuing-attitude is not held by the addressee – if condition (ii) is not met – then the request fails entirely and does not create a reason at all. And this outcome is one that might seem problematic. That is, one might think that even when a request is made by a contemptuous fiend – a person toward whom one holds no discretionary valuing-attitude whatsoever – a reason might nonetheless be *created*.[69] The thought is that the reason

---

[68] A question that might be raised to my view (and indeed has been, by Alfred Archer, to whom I am duly grateful), is whether the reasons that are thus created by requests are free-standing considerations that favour Φ-ing, or whether they can only ever play an accompanying role to other reasons that must also be at play. Specifically, does the reason that is created by the request (qua request) depend on there being a favouring reason that stems simply from the existence of the addressor's desire for the addressee to Φ. In the terms of Jonathan Dancy's work on the different sorts of practical reasons that there are, this is the question of whether requests create 'favourers' or 'intensifiers' (Dancy 2004, pp. 38-43). My answer is that the reasons created by requests are stand-alone favourers. When we hold an interpersonal valuing-attitude in another we endow another with the power to create reasons by requesting. Conceptually speaking, I do not see why it should be impossible to endow someone with this power whilst for whatever reason, not treating their very desires as themselves the sources of practical reasons.

[69] I am very grateful to Glenda Satne for pushing me on this point.

might be created just as the fiend intends even though in the addressee's deliberations it has insignificant normative weight, or is dramatically outweighed by countervailing considerations against acting *for* the contemptuous addressor.

But this problem does not arise if one keeps in mind the structure of requesting as a distinctive normative operation – a structure that depends on appropriate mutual acknowledgement of the intention to create a reason. Recall that in the theory of robust reason-giving, to which my proposal is an amendment, requests create reasons only when the addressor intends the addressee's given reason to φ to depend in an appropriate way on the addressee's recognition of the addressor's communicated intention (to give the addressee a reason to φ). On my proposal, the addressor intends to trigger a reason which itself depends on the addressee's discretionary value-outlook. As such, the addressor intends the request as an appeal to an item in the addressee's own discretionary value-outlook. If there is no such item, if the addressee does not place any discretionary value in the person of the addressor, then the request fails to create a reason. Moreover, it fails to create a reason *even by the addressor's own lights*. Of course, the fiend may succeed in coercing the addressee, or the fiend may have the authority to command her, or it is even possible that the contemptuous fiend can reveal his desire (epistemically) for the addressee to do his bidding, and that mere desire may give the addressee a reason. The fiend may be successful in creating reasons in all of these sorts of ways. But without the addressee holding discretionary value in their addressor, the latter cannot create a reason in the normatively distinct sense of requesting.

As a further illustration of the theoretical worth of the proposal being made here, consider the case of entreaties. I shall use the term 'entreaties' to refer to a subset of requests in which, prior to the request being made, condition (ii) is not met: the addressee does not yet place sufficient discretionary value in their addressor to grant them the standing to make the request.[70] In entreating, one attempts to trigger *both* conditionals of the prior reason. That is, the addressor appeals to their addressee to actively place discretionary value in them – in the person of the addressor – and simultaneously to request, on the basis of that discretionary valuing-attitude, that the addressee undertakes some action.

Suppose that Carrie and Anita are strangers to one another. Carrie is walking down the street on which Anita lives and she urgently wants somewhere to hide, but to tell anyone why she needs to hide would risk endangering her confidant. She knocks on the door of a house on the street and Anita answers. Carrie asks whether she can come in, without offering any explanation. We might imagine that she asks whilst looking directly into Anita's eyes. In the moment before this exchange, if Anita had been asked 'would you consider the request of a stranger to come into your house, without explanation', she would have said no. Anita would not have granted a stranger even the standing to make that request – not without some explanation. But in the moment of the entreaty, Carrie

_____

[70] In giving this specific meaning to the term entreaties, I am following Lance and Kukla (2013, p. 474): 'an entreaty is a meta-call: it calls someone to grant the caller an entitlement to make certain kinds of claims that the caller is not yet in a position to make.'

implicitly appeals to Anita to take up some kind of valuing-attitude toward her. This could be admiration, affection, pity, some kind of endearment. (As I mentioned in section 2 above, these particular feelings are only points in an indefinite range of favourable interpersonal attitudes that one can hold, or not, at one's discretion.[71]) It is a presupposition of Carrie's entreaty that the following conditional reason is true: Anita has reason to let Carrie into her house, if Carrie asks, and if Anita places sufficient discretionary value in Carrie.

A theoretical advantage of my proposal is that it equips the theory of robust reason-giving with the capacity to explain the normativity of exchanges like Carrie and Anita's, of entreaties in general. This advantage is pertinent because, as it seems to me, the boundary between entreaties and ordinary, run-of-the-mill requests is fluid and often difficult to identify. A run-of-the-mill request, let us say, is one in which condition (ii) is met already, prior to the request being made. All it takes is for the request to be uttered and a reason will have been created for the addressee, without any alteration in anybody's value-outlook also being required. For example, Gwen and her younger sister Cat are being looked after by their babysitter, Wynn. Gwen is extremely enamoured of Wynn – thinks the world of her – so when Wynn asks Gwen to go and read Cat a bedtime story, there really is no question of whether she places sufficient discretionary value in Wynn for the request to be reason-giving. This is a clear-cut case of a run-of-the-mill request. But I suggest that often, depending on the demandingness of the action that is being asked for, things are less clear. When we make requests, it seems that we often implicitly appeal to our addressees to value us – we seek to convey our worthiness of pity or esteem, or whatever, as a way of bolstering the reason that our request attempts to provide. On my view, these appeals may sometimes be requirements for the request to succeed in creating a reason at all. If the border between entreaties and run-of-the-mill requests is indeed as hazy as this, then any account of requests should be capable of explaining *at least* how there could be such a hazy border. Enoch's account cannot, since the interpersonal valuing attitudes that are appealed to in entreaties play no role in his theory of requests. But my proposal explains the possibility of the hazy border, as well as giving an account of the normativity of the terrain on both sides of that border: of both entreaties and run-of-the-mill requests.

In this section I have presented an amendment to the theory of robust reason-giving. The amendment helps to explain the sense in which the reasons presented in requests are discretionary reasons. I will now raise an objection to my proposal in the form of a competing explanation of the discretionary quality of requests, one that is popular in the way that speech act theories think about the matter. I will argue that this competing explanation fails to address the normative questions that my view sets out to confront.

---

[71] Having said that, Cristina Roadevin has pointed out to me that some requests make their appeals to *particular* interpersonal valuing attitudes. For instance, the lover's request is appropriately granted out of love, affection, but not out of, say, pity.

5. The pragmatics of discretion[72]

A central merit of the account that I am proposing is its capacity to explain the sense in which the reasons presented in requests are discretionary. It does so by appealing to the normative role played by discretionary interpersonal valuing attitudes in grounding those reasons. As such, one way to challenge this account would be to offer an explanation of that discretionary quality without recourse to such interpersonal valuing attitudes. A rival explanation of exactly this sort is to be found in the approaches by speech act theorists to the phenomena of requesting.

Such approaches tend to follow John Searle (Searle 1969, p. 62) in thinking of requesting as expressing a desire of the addressor for the addressee to undertake an action. As such, these approaches concern themselves with what I earlier characterised as a form of epistemic reason-giving. They do not address what I argued in section 1 to be the distinctive normative role of requests as such, wherein a request does something more than merely convey information (either about the addressor's desires, or anything else). But what is more interesting for my purposes is the way in which such pragmaticists think about the discretionary quality of requests in terms of the *indirectness* of requests as speech acts (Searle 1975). This analytical perspective distinguishes between degrees of (in)directness (Kádár & Haugh 2013, pp. 23-25). Thus, while a straightforward imperative might be possible, an addressor has the option of deploying layers of indirectness. This could be achieved by phrasing the request as a question ('would you please…'), adding qualifications ('if you wouldn't mind…'), or even merely implying the request by making a related assertion (as when the assertion, 'It's a little cold in here,' implies the request to close the window). In some instances, an addressor may choose to make a request out of politeness to the addressee. In other instances, this indirectness may function to protect the addressor themselves against embarrassment in the event of the request being refused. Either way, the indirectness of the request is a mechanism by which the addressor communicates their acknowledgement that the reason that they are presenting is discretionary. This is the important point. By focusing on the mechanics of indirect speech acts, one can conceive of the discretionary quality as simply this: an acknowledgement by the addressor, concomitant with the request, that the request creates a reason that the addressee could heed or not, at their discretion.

To be sure, this perspective does look like a challenge to the proposal that I am advancing. The challenge holds that all there is to the discretionary quality of a request is explicable in terms of the communicated acknowledgement of the optional or discretionary force of the reason. This rival explanation threatens to make the idea of the discretionary value of persons superfluous to a theory of requests.

But I do not think that what we have here really is a rival explanation. I do not contest that requests can be made with varying degrees of indirectness. Nor do I contest that such indirectness can, to varying degrees, indicate the addressor's willingness or preparedness to accept the refusal of the

---

[72] I am indebted to Basil Vassilicos for raising this challenge and helping me to think it through.

request. But the question that I have been addressing in this chapter is what normative force, if any, there could possibly be to the discretionary reasons that we create in requests. And to this, as far as I can see, the pragmatic analysis does not propose an answer. In fact, therefore, the pragmatic analysis of discretion begs precisely the question that motivates my proposal.

To see the point here, it may be helpful to attend to the contrast between the kind of reason one can intentionally give another by virtue of being authoritative (a command), and the kind one can give without authority (a request). How could the former kind of reason be discretionary? It is difficult to see how a reason could at the same time derive its normative pull from the authority of the addressor, and be discretionary in the sense of decidedly *not obliging* the addressee. Reasons stemming from addressees' authority are – surprisingly enough – authoritative, they are commanding, they are non-obligatory. Having the authority to command someone to φ does not necessarily entail having the power to give them reasons to φ with watered-down, less than obligatory strength. As such, the pragmatic analysis, in illustrating the mechanisms through which people present discretionary reasons, thereby illustrates the existence of interpersonally-given reasons that do not derive from the authority of the reason-giver. That is the phenomenon targeted by my suggestion for a theory of requests.

Conclusion

I began here by stipulatively defining a request as an attempt by an addressor to create and communicate a non-obligatory reason for the addressee(s) to perform an action. Beyond merely stipulating this definition, I have argued that it captures the conceptual heart of the notion of request, even though it is also fails to match up with the breadth that the term takes on in ordinary usage. On the basis of this definition, I have motivated a general philosophical question about this aspect of our practical lives: what could possibly be the normative grounds of such reasons as those created by requests? In light of the discussion of Enoch's theory of robust reason-giving, of my own proposal, and related objections, I now have an answer to that question. The normative ground of a successful request, and the reason that the addressee has to accede, is the truth of a prior reason with two conditional elements. That prior conditional reason is of the general form: Person A has a reason to φ if requested to do so by Person B and if Person A places sufficient discretionary value in Person B. As such, the reasons presented in requests make essential reference to the discretionary value outlooks of their addressees. In particular, they rely on the addressee holding a certain degree of interpersonal valuing-attitude in the person of the addressor. The normativity of requesting, therefore, is a product of a deeper normativity: that value that people have for one another, which can wax and wane in the course of interpersonal interaction.

Chapter 5
Love first, ask questions later

It arose from my introductory discussion in Chapter One, that there are a number of different conceptions of what it is to relate to another second-personally. In Chapter Three, I singled out one such conception in particular, namely that A's relation to B is second-personal if and only if A recognises B as a partner in intentional interaction, such as address. The very idea of a contrast between a stance of addressing, on the one hand, and a stance of regarding things in a non-addressing manner, on the other, is a distinction that was brought to prominence by the Austrian-born philosopher and theologian Martin Buber. Buber himself did not think of the second-person relation in quite the terms that I have used so far, but his concept of the I-Thou relation is an ancestor, to which my concept of the second-person relation bears a notable resemblance. In this chapter, I will turn to Buber's thought directly in order to investigate a further aspect of the practical significance of the second-person relation.

In Chapter Two, I advanced a sceptical view of Darwall's attempt to show that moral obligations are dependent for their normative force on second-person relations. However, in this and the next chapter, I will offer alternative angles from which the second-person relation might play a role in the normativity of moral reasons. In Chapter Six, I will consider whether the second-person relation can alone provide epistemic access to the value of persons – which is a factor in determining swathes of moral reasons, including those at play in moral obligations. First though, I will examine the second-personal character of agapic love, which in at least some contexts seems to give rise to moral reasons. It may be that there are many circumstances in which one could act lovingly towards another person, but where to do so is not morally obligatory. To love the other in such circumstances would be a morally good thing to do, but would go beyond what one's duties require. Such love is supererogatory. It is especially because of its promising connection with morality that the second-personality of agapic love is of concern to this dissertation.

Agape is the love of people as people. It is a kind of loving-attitude that anyone can take towards anyone else: towards total strangers, and sworn enemies; towards those who seem for all the world to be utterly unworthy of love. And in being a loving-attitude, agape is a relation towards a person in which one can see and be moved by reasons to be kind to them, and in which one is emotionally committed to the value of their life and their ends. A puzzle though, is whether this agapic relation towards others is simply a blind, unconditional projection of warm-heartedness, or a rational response to the way they are. That is, is agape a relation characterised by openness, or by attention to the other? In this essay, I will consider these two options. After presenting their strengths and the tension between them, I will offer a sketch of how, despite appearances, openness and

attention might form compatible and indeed complimentary attitudes that together comprise the agapic relation, where that is an interesting candidate for the morally best way to relate to another as a person.

In order to see the importance of gaining a theoretical understanding of agape in particular, one must see how it is distinct from other kinds of love. I propose to construe agape as distinguished by the breadth of its appropriateness. Agape is like other kinds of love – romantic love, friendship, familial love, for instance – in being a bundle of warm-hearted dispositions of one person towards another. All kinds of love consist, perhaps *inter alia*, in dispositions to care about, feel certain emotions for, and act in the interests of, the beloved. But whereas romantic love, friendship, and familial love (as well as other types) are each appropriate only within the context of particular relationships and interpersonal situations, agape is putatively appropriate in just about any interpersonal situation. So agape is just that kind of love which is appropriately felt by anyone towards anyone else, if it is appropriate at all.

In being more generally appropriate than other species of love, agape can be distinguished from them by its lack of each of their own defining characteristics. Unlike romantic love, agape does not necessarily include any dispositions to desire intimacy with the beloved, sexual or otherwise. Unlike friendship, agape need not involve a desire for friendly interaction, or for the subject to be an integral part of the life of the other. Unlike familial love, of course, agape does not require the lover to see themselves as part of the same exclusive social group as the person they love. Indeed, it is most famously thought of as a kind of neighbourly love that is exemplified in love of 'Others' (with a capital 'O'), who one sees as outside the community to which the lover themselves belongs.

In its paradigmatic form, agape is thought of as morally praiseworthy. Other kinds of love might also be morally good, but agape's purported general appropriateness in interpersonal relations marks it out as a candidate for the morally best way to relate to other people in general. There is some forceful intuitive appeal to this thought that it is appropriate and indeed morally good to love others in general, including those who seem unworthy of such love. Whether in fact it is generally appropriate to love others is an important normative question. A good answer would consider the detail of apparent objections, like the moral case for shunning or even being actively hostile towards flagrant wrongdoers, towards those who are resolutely spiteful, manipulative and cruel.[73] I am not going to consider these issues, however, and the normative question is not my question in this chapter. My task, rather, is further upstream. It is to articulate what kind of relation agape must be, if it could plausibly do justice to the sentiment that it is morally good to love others in general.

---

[73] For a forceful moral critique of the notion of agape, see (Andolsen 1981)

1. Openness

Martin Buber (1958) is best known for his notion of the *I-Thou* relation. This is the idea that there is a way for a person to relate to another person – as an I to a Thou – where that other ceases to be regarded in the way that the subject customarily regards things outside of itself, as just another object to be observed, comprehended and (where possible) used as a means to the subject's ends. Rather, in I-Thou relations, the Thou is not regarded as an object, but simply as the opposite pole of a relation of dialogue with the subject:

> When *Thou* is spoken the speaker has no thing. He has indeed nothing. But he takes his stand in relation. (1958, p. 4)

Love, for Buber, is native to such I-Thou relations, where the other is not judged on the basis of their objective properties. As such, Buber does not regard love as a sentiment that is in the first instance based upon an appraisal of the other's *worthiness*:

> In the eyes of him who takes his stand in love, and gazes out of it, men are cut free from their entanglement in bustling activity. Good people and evil, wise and foolish, beautiful and ugly, become successively real to him; that is, set free in their singleness they step forth and confront him as *Thou*. (1958, p. 15)

Here is a conception of love where characteristics of a person that might otherwise be thought to merit taking positive valuing-attitudes towards them are drained of any significance at all. Thus, one aspect of Buber's thought might be formulated in the following thesis:

Agape as Openness   For a subject, S, to agapically love another, X, S must disregard all properties of X that may render X worthy or un-worthy of love.

This is definitely only an aspect of Buber's understanding of the openness of an I-Thou relation. He goes much further than this thesis, claiming that relating to another as Thou requires disregarding absolutely all of their qualities including any identifying characteristic; including, for instance, their spatiotemporal location. Moreover, he explicitly rejects the idea that love could be comprised of the way that parties to a love relation relate to one another, where such 'relating' amounts to having certain mental states with the loved-one as their object (1958, pp. 14-15). That is, he does not think it possible to give an account of love by explaining the kinds of beliefs, feelings, attitudes and dispositions that are features of individuals' experiences of love. He thinks that the phenomenon of love is something that lies outside of individual mental states, and so is '*between* I and Thou' (1958, p. 15).

Without wanting to dismiss the possibility that there might be something worth exploring in that claim, I will put it to one side for present purposes. There is a thought lurking between the lines in Buber's text that is attractive precisely because it seems to capture something of what it is like to love others. And thus, insofar as it rings true to an experience, the thought must be about the mental states involved. So at this point I shall depart from the letter of Buber's text (but, you understand, in the hope of getter closer to its spirit). To be explicit, the thought of Buber's that I find alluring here is that

some of the time, loving others involves having loving attitudes towards them irrespective of anything about them that might qualify them (or disqualify them) for such love. The object of such attitudes is the other themselves, rather than them under some description according to which they have the kind of properties that could be love-worthy. And what makes these passages from Buber relevant to an *enticing* conception of agape, is that this kind of openness towards others plausibly figures in a model of how we ought to relate to others just as people – including to strangers and those we are inclined to dislike.

There are two points to make in support of Agape as Openness (hereafter just *Openness*). First that it is sometimes *right* to relate to others in such a way as to abstain from evaluating their character, but to love them anyway. And second, that such non-evaluating love is not only possible, but is a recognisable feature of social life. Each of these points can best be illustrated through examples.

Consider the Biblical expression of agape (King James Version 2011):

But I say unto you, Love your enemies, bless them that curse you, do good to them that hate you, and pray for them which despitefully use you, and persecute you;
– Matthew 5:44

The proper way of relating to people, Christ says, is to love them. Love other people, and extend one's love even to those that curse and hate you. This can seem quite a challenging pronouncement. Indeed, at first blush, one can readily imagine circumstances in which loving a particular person seems to be a distinctly *in*appropriate response to their character, or to their behaviour; circumstances where a more fitting response would be something closer to contempt than to love. The attraction of the Buber-inspired *Openness* is that it can make sense of something like this Christian ideal even in those challenging circumstances.

As an example, consider Arwen's attitude to her colleague, Dilys. The two have worked in the same office for a few years, during which time Arwen has never really warmed to Dilys. In fact, while her first impression was of a slightly ostentatious, slightly snobbish newcomer, that impression has only worsened as further aspects of Dilys' character have become apparent: she's selfish and proud, and quite brazenly sneers at her colleagues. In short, Arwen has never liked her. So now imagine an episode in this relationship when two events come in quick succession. First, Arwen gathers (perhaps from a video taken on a staff night out and shared privately with Arwen online) that Dilys has been speaking unkindly of Arwen behind her back: making jokes at her expense, casting aspersions about her private life and generally speaking about her with bitterness. Arwen, naturally, is upset. But that's not all. Later on the very morning of this revelation, Dilys comes to Arwen's desk with the office diary in her hand, and asks Arwen whether she'd mind swapping their days-off in a fortnight's time.

Suppose that from the perspective of Arwyn's self-interest, it is neither here nor there whether she accedes to the request. In that case, does she have reason to do Dilys this favour, to agree to swap their days-off? Or, in more general terms, do we have reason to perform supererogatory altruistic

deeds in the interests of others who we take to be unworthy of such kindness? Dilys' offences might seem relatively trifling, but this is a case where Arwen cannot see any aspect of her character that calls for kindness, that would justify love. And yet – tryingly – this case is clearly one to which the Biblical prescription of agapic love is applicable.

Openness offers a way to make sense of this prescription. It says that what one ought to do is relate to the other themselves, shorn of all evaluable qualities, and love them. That is to say, without judging them at all, one should be disposed to feel concern for them, and to do them good. As such, *Openness* proposes that one disengage one's rational faculties from one's heart, so to speak. In this instance, in order to be moved by reason to act with loving kindness towards Dilys, Arwen ought to relate to Dilys as a *Thou* in Buber's sense.

The idea of loving someone without judging them is not as mysterious as it might sound. It is just the idea that what might otherwise seem to be properties that should figure in whether one sees someone as love-worthy, do not in fact figure. In the example given, one can imagine Arwen summoning some resolve in the face of Dilys' hurtful behaviour, and subsequent request. One can then imagine Arwen, from a certain state of mind, reflecting on Dilys' character sympathetically, in such a way that the sheer independence of her interior world from Arwen's own, could seem strikingly salient. From this perspective, those traits that moments before had seemed to warrant a sentiment of quite some antipathy – Dilys' sneers, her pride, and her misdemeanours – could suddenly come to seem insignificant. Perhaps Arwen could re-envisage these as reflections of Dilys' own frailty; or perhaps she could simply come to see them as such petty matters as to not bear any weight. In any case, the kind of perspective that Arwen could take towards Dilys is one where any consideration of particular facts about her or about her deeds, carries no relevance in determining Arwen's attitude towards Dilys herself, as the person whose own life and ends is far greater than, and separate from, any summary of such facts. From this perspective, Arwen illustrates how one can love another as a Thou.

Importantly, such openness as would be required of Arwen were she to succeed in acting lovingly towards Dilys in this instance is not a miraculous or even unusual attitude. This is the second point mentioned above in favour of *Openness*: that it is a recognisable feature of social life. Indeed, whenever one person takes a strong dislike to another – or hates or is repulsed by them – such antipathy is not necessarily permanent. That is, sometimes these negative interpersonal sentiments can soften and melt away. And when we reflect upon those instances – occasions of dislike shading into some more positive feeling – an attitude like openness may often be apparent as a crucial transitory stage in the process.

Consider, for example, Catrin's relationship with one of her dance students, Rhian. Since their first lesson, Catrin has been increasingly peeved by Rhian's petulance. Every attempt of hers to give Rhian some instruction or criticism is met with defensiveness, which offends Catrin and typically prompts further criticism. Other students in the class are neglected as Catrin's energies – against her

better judgement – are absorbed in her somewhat combative relationship with Rhian. Catrin is eminently professional and endeavours to treat Rhian just as she would any other student. But sight of Rhian does not prompt feelings that could be said to resemble love. In fact, in Catrin's eyes, Rhian's discernible characteristics all offer evidence of her bratty character: the shrill pitch of her voice, her sprightly gait, her easy chattiness with her classmates.

So imagine an encounter between the two in which Catrin's contempt for Rhian begins to soften. Suppose that for some extraneous reason Catrin is in a particularly good mood, and perhaps she finds herself choosing Rhian as a partner with whom to demonstrate a new routine to the rest of the class. As the two move through the steps, Rhian's characteristic ostentatious stubbornness is made manifest, but for once Catrin is not annoyed. She feels warmth towards Rhian and enjoys the interplay of their bodies before the class. She surprises herself by indulging in joining Rhian's performance, rather than resisting it. In this moment it is not that those traits in Rhian that previously inspired Catrin's irritation are no longer visible, but rather that they just don't seem to matter. Plausibly, in cases like these, pejorative ways of seeing others can be *drained of their significance*. That is to say, the stance of openness that one person can take towards another is one that facilitates love by neutralising one's evaluation of the other. It facilitates love without needing to rationally undermine negative judgements – i.e. by showing the judgement that Rhian is a brat to be mistaken – and also without needing to provide the basis for a countervailing positive judgement – i.e. by revealing Rhian's underlying character to be ultimately endearing.

The stance that Catrin takes towards Rhian on this particular day is exemplary of the I-Thou relation in the sense that is relevant to present concerns. Catrin relates to Rhian as a Thou, in that she sees Rhian as a person, as a partner in a distinctive kind of interaction and in that moment thereby transcends Catrin's own evaluative judgements of Rhian that would normally filter her engagement with her. And, significantly, this interaction illustrates that instances in which people take stances of openness towards others could be placed on a spectrum: from saintly and indeed rare moments of (for example) profound unconditional forgiveness, to the rather more prosaic occasions in which minor gripes slip away and disappear.

To recap, I have suggested in this section that Martin Buber's notion of an I-Thou relation provides a model for a way of thinking about agapic love. I have further made a case that the kind of openness embodied in an I-Thou relation is a morally desirable attitude to take up, since it may enable one to see and be moved by supererogatory altruistic reasons from which one is otherwise alienated. Moreover, I have illustrated that this openness is a realistic attitude: beyond being merely possible, it is a common moment in many of our relationships.

2. Attention

Simone Weil articulates a contrasting notion of agape. A core concept that Weil develops in several parts of her work is that of *attention*. Attention refers to a kind of mental action where a subject deliberately comports themselves towards something with a view to becoming acquainted with the good (or goodness) that resides in that thing. Paradigmatically, attention is something exercised in interpersonal relations. The way that this contrasts – starkly – with *Openness*, is that the purpose of this mental activity involved in attention is to gain a closer, more detailed appreciation of the characteristics of the object to which one is attending, and the ways in which those characteristics are good. Here, agapic sentiment arises precisely by directly considering the other's love-worthy properties, through as sophisticated and patient an evaluative effort as one can muster. Consider this passage, where Weil brings the notion of attention specifically to the context of interpersonal relations.

> [I]t is indispensable, to know how to look at [our neighbour] in a certain way… This way of looking is first of all attentive. The soul empties itself of all its own contents in order to receive into itself the being it is looking at, just as he is, in all his truth. (Weil 2009, p. 115)

Attention is a matter of truth seeking. The visual metaphors of looking and seeing are a recurring device in Weil's attempts to convey the nature of this attitude: it is one where the subject's cognitive capacities are engaged and active, while the subject's evaluative prejudices, and the nag of their self-interest, are – as far as possible – neutralised. That is what we should understand by the soul 'emptying' itself of its own contents. But attention is not just the attitude taken up *whenever* cognitive attitudes are engaged. Rather, it is a technical term in Weil's vocabulary specifically referring to the procedure of coming into acquaintance with the good via an appreciation of the specific qualities exhibited by good things. She says:

> If we turn our mind towards the good, it is impossible that little by little the whole soul will not be attracted thereto in spite of itself. (Weil 1952, p. 106)

This emphasises that the kind of acquaintance that a subject gains by attending to the good is not merely the awareness of a spectator – attending to the good inevitably leads to being *invested* in it, to actively valuing that which manifests the good. In the context of the present discussion, attending to another and becoming acquainted with the goodness in them just is coming to value that person, lovingly, for that about them that makes them worthy of such love.

In her own philosophical outlook, there is a religious significance to this possibility of meeting goodness itself through rich contact with the detail to be found in worldly things. Presently I will try to show that the essence of Weil's concept of attention is an appealing way of describing an important aspect of interpersonal love, even in a completely secular worldview. First, here is the conception of agape that could be drawn out of Weil's notion.

Agape as Attention   For a subject, S, to agapically love another, X, S must attend to X's love-worthy properties.

This is a way of understanding agape that is diametrically opposed to the Buber-inspired *Openness*. The latter proposes that to love a person agapically one ought to relate to them in such a way that their individual characteristics, all their evaluable qualities, are drained of significance, and one ought to love them irrespective of any such qualities; whereas this counter-proposal drawn from Weil replies by asserting the role of such individual traits of others – and the minutiae of things in general – in grounding agapic love.

Famously, this insight of Weil's was taken up and developed by Iris Murdoch in the first essay of *The Sovereignty of the Good* (Murdoch 1970, pp. 16-7) where Murdoch elaborates on the notion of loving attention through an example of a woman and her daughter in-law. The gist of the example is that the woman came, over a long period of solitary reflection, to alter her attitude towards her daughter in-law. Having once thought of her in distinctly unfavourable terms, the process of reflection led her – painstakingly – to recognise the attractive features of her daughter in-law's character. Through this process, the pejorative descriptions of her daughter in-law dissolved away and left the woman with a genuine affection for her, based on her actual, underlying worthiness for that affection. In Agape as Attention (*Attention*), I mean here to appeal to the same vision of love that inspired Murdoch.

There are several attractions of *Attention* as a way of spelling out the general stance that it is appropriate to take towards people (as people). Most distinctively, it captures an important point about our *fallibility* in judging one another. That is, the aspects of people that reveal (or perhaps comprise) their true worthiness for love can all too often be difficult to see on a first appraisal. Indeed, perhaps this is a consequence of the pace of modern life, of the economic grind within which our interpersonal interactions play out: but it may likely be the exception, rather than the norm, to be attuned in one's relations with others to their love-worthy qualities.[74] Weil's insight is that contrary to this custom we ought to focus our attention in our relations with others on that about them that makes them worthy of love, even if those qualities lie beyond initial appearances. On epistemic grounds, as well as moral ones, it seems right that if our ordinary attitudes in our interpersonal relations blind us to certain facts – in this case facts about others' love-worthy qualities – then those ordinary attitudes are quite clearly improper. For Weil, turning our attention deliberately towards others in the right way will lead inevitably to an acquaintance with the good that resides in them. As she puts it in the quote above, 'the whole soul will be attracted thereto in spite of itself'.

Accordingly, one might think that Arwen should not be so sure of her harsh judgement of Dilys. Perhaps her unfavourable first impression was borne of a misunderstanding, and perhaps that initial tint of suspicion subsequently came between the two, preventing Arwen from seeing Dilys' actual character, replete as it may be with virtues and charms. This epistemic worry motivates the

---

[74] A similar concern – that modern life typically involves seeing others in an *impersonal* light – is what agitated Adorno (2005, p. 23) when he said that 'the entire private domain is being engulfed by a mysterious activity that bears all the features of commercial life without there being actually any business to transact'.

thought that the proper – morally best – way for Arwen to relate to Dilys would be to make a conscious effort to see Dilys in a more favourable light in the hope of thereby becoming acquainted with love-worthy attributes of hers, which may currently be obscured from view.

There is some intuitive pull, then, to thinking of this kind of loving attention as a feature of agapic love. But this is far from the only advantage of *Attention*. As an account of how we ought to relate to others generally, it has the further advantage – which *Openness* lacks – of conforming to a widely held rationalistic conception of love. Let's ascribe the label 'Rationalism' to the view that love necessarily involves – among other things – the belief that the object of love is worthy of love, all things considered.[75] Figures in the recent philosophical debates about love differ widely on what sort of reasons could justify love, but there is a notable degree of convergence of opinion on Rationalism.[76] Since their respective in-/compatibility with Rationalism is an important distinction between *Openness*, and *Attention*, it will be worth briefly setting out the core motivations behind that Rationalist conception of love in general, before reviewing the standoff between the two putative accounts of agape that I have articulated.

Among the reasons to think of love as a rational attitude (in a broad sense, as already specified), three considerations in particular stand out. First, and most basically, Rationalism explains the sense of ownership and responsibility that one has for one's loving attitudes, which is much like the ownership and responsibility that one has for one's beliefs. The process of coming to love someone is often not a voluntary process: in romance, there is little choice involved in falling in love, and the same can usually be said about the formation of other loving bonds, with family and friends, for example. But despite love's involuntary origins, it is an attitude (or family of attitudes) for which agents are responsible, and with which, typically, they identify themselves. In this regard, loving attitudes differ from other involuntary features of our behaviour. I am inclined to defend the appropriateness of my love of my friend, but not the appropriateness of my coughs, sneezes and hiccoughs, for instance.

A second, closely related motivation for Rationalism is that it explains the possibility of *inapt* love. Only if love necessarily involves a cognitive component – cashed out as a belief in the all-

[75] This terminology is in-keeping with that used by Jollimore (2011, esp. chpt. 1) in his influential book on love.
[76] Some figures, including Jollimore (2011; 2015), Abramson & Leite (2011)and Protasi (2016) defend the view that facts that mark a person apart from others can at least partially justify loving them – that is, things like their beauty, kindness, humour, and so on. Others, like Kolodny (2003) and Rorty (1987) think that historical facts about the relationship between lover and beloved can justify the former's love. Still others, like Kreft (under review), Setiya (2014), and Velleman (1999), think that it is the bare personhood, or humanity, of the beloved that actually justifies the love, even if it was more particular, individual traits that caused the attraction from which the love arose. Without discussing any of these schools of opinion directly, the view I advance in this chapter tries to stay neutral between the first and second groups – on the matter of exactly which kinds of individual properties of a person and / or her relationships with others can figure in a justification of loving that person. My view is also offered as an incompatible rival to the bare-personhood views, since unlike them, I argue that love must pay heed to individual features of the beloved. For present purposes, what is most noteworthy is that despite such a diverse debate about the nature of the justification of love, there is wide agreement on the fact that love is the sort of thing for which justification could (and therefore should) be given. There are, nonetheless, some dissenting voices even to this widely-held view, cf. Frankfurt (2004).

things-considered love-worthiness of the loved-object – can one make sense of love being the wrong attitude to take towards someone in some contexts. It is irrational to romantically love people who, as a matter of routine, have only ever spitefully abused your lovingness. Indeed, it is also irrational – and therefore wrong – to treat such people as friends. Another way of stating this motivation is really just that love seems to require some kind of justification.

This leads to the third core motivation for Rationalism, which is simply that we are often conscious of reasons that justify our loving attitudes. If those attitudes were not in principle cognitive, then justifications would not be relevant or required. And yet we often are able, when asked, to answer the question 'why do you love them', with non-instrumental reasons: they are charming, virtuous, we depend on them, and they on us, and so on.

A popular line of thought would have it that whereas Rationalism holds true for some kinds of *personal* love (friendship and romantic love in particular), agape, by contrast, is non-rational.[77] This line of thought renders agape as a saintly stance, one quite removed from our ordinary interactions. I think this is a distinct misunderstanding. Paradigmatic instances of agapic love are between neighbours, fellow citizens, strangers, members of conflicting clans or sects etc. These are the relationships that most readily spring to mind as examples of people who are not acquainted with one another's unique love-worthy properties. But to suppose that loving attitudes in such relationships as these cannot be rational would be a mistake. Nothing precludes strangers or enemies from *becoming* acquainted with one another to the extent that they come into contact with, and are moved by, one another's worthiness for love. To become acquainted with another person and to come to love them despite one's initial isolation from their merits is part of agapic love, and it is decidedly rational. Consider someone who encounters a stranger, feels uncertainty and suspicion towards them, but moves beyond these inhibitions and comes to appreciate the stranger's charms. This person's newfound love of the stranger is rational. It implies a belief in their all-things-considered worthiness for love. The agapic lover is the owner of her love and is responsible for it. It could be inapt – the stranger's properties could, when all things are considered, speak against loving them. But when it is apt, the agapic lover has a sense of what makes it apt.[78] So at least insofar as Weil is right about the kind of love that ought to be directed towards one's neighbours, towards people in general, agape is no less of a rational attitude than the other species of love.

By contrast, in eschewing judgements of love-worthiness, *Openness* is at odds with Rationalism. If to love someone agapically is to refrain from judging them, but to love them regardless, then the notion of love at play is, from a Rationalist perspective, somewhat suspect. Whatever such loving amounts to, it seems that it may lack several qualities that are often thought of

---

[77] For instance, here's Jollimore (2011, p. 176 n.2): 'I would... place the Christian idea of agape in a category separate from friendship and romantic love. Agape might represent or involve any number of things – beneficence, goodwill, charity, etc. – but it involves no appreciation or attention directed towards its object.'
[78] Though this is not to say that reasons for love are always perfectly transparent. Indeed, they are often opaque: we are not always able to say what makes our beloved quite so lovable.

as being characteristic of love generally: being an attitude for which agents take axiological ownership, an attitude which in some circumstances could be inapt, and being such that those who hold the attitude can be conscious of reasons that justify it.

There is thus a deep tension between these two views of the nature of agape. *Openness* is a non-rational love that transcends discrimination. It is a love of others *as themselves*, in some sense, shorn of their contingent, material characteristics. In being indiscriminate, this view of agape is closer to the traditional Christian notion, and closer to some canonical philosophical articulations of that notion, such as those in Nygren (1982) and Kierkegaard (2013).[79] On the other hand, the Weil-inspired *Attention* is revisionist of that traditional notion. But as such, it regards agape properly as a kind of loving attitude – as having a closer affinity with other species of love than is traditionally assumed.

Whilst I hope to have drawn out the intuitive appeal of these two attempts to make sense of the love that we all ought to bestow upon others as people, I have not offered any defence of the supposition that the morally best way to relate to others is a relation of love at all. It should remain a live possibility that the idea of universal love belongs to a religious worldview, but not to a secular one; that a fully-fledged philosophical picture of the basic, ideal interpersonal stance will find something like respect, rather than love, to be the standardly appropriate attitude. But attempting to spell out a theoretical account of the normative grounds of agape must remain a separate project. The remainder of this chapter will pursue my more modest, preliminary goal: to articulate a coherent story of what the attitude of agape could consist in, in view of the antagonism between the two views unearthed above.

3. Retroactive justification

The two accounts of agape that I have spelled out both have their appeal. But they are contradictory. *Openness* says that agape is a stance of paying no heed to the specific properties a person bears; *Attention*, that it is precisely one of paying heed to such properties. This contradiction is not an artificial construction, but a genuine opposition between intellectual perspectives on agapic love. The debate about the rationality of love has focused heavily on personal loves – of friends, or romantic love – and that debate is shaped by a disagreement. There are those who think that people's individual characteristics play a role in justifying loving them, and those who do not. The same disagreement extends to agape. Those who deny the role of individual characteristics in grounding personal love would deny that they play any role in grounding agape either.[80] Proponents of such a view may therefore sympathise with *Openness*. Contrastingly, those who affirm the grounding role of individual

---

[79] For a survey of the traditional conceptions of agape, see (Outka 1977).
[80] Velleman's (1999) and Setiya's (2014) accounts, for example, are already geared-up to extend their scope and explain the rationality of agape, where they would see agape as being grounded in the other's 'bare-personhood' or 'humanity', respectively, *not* in the other's individual characteristics.

characteristics in personal love would construe *Attention* as really just a particular strain of personal love, but they would consider *Openness* to be something altogether different. So the contemporary debate about personal love reflects the significance of the opposition that I have drawn out between the two – rational and non-rational – conceptions of agape.

Despite the fact, then, that the tension between Buber and Weil is enduring, and significant, I think it is misleading. Agape is best understood, I suggest, as a loving stance comprising a series of two episodes – an episode of openness, and one of attention. If agape were to be justified as a rational stance to take towards others in general, then it must amount to the following doctrine: of any particular other person, one has defeasible reason to believe that they embody some property, X, in virtue of which one would be justified in loving them.[81] The love bestowed on the basis of this doctrine lacks an acquaintance with that in the other which is supposed to justify the love. It is an openness which calls for attention.

Earlier, I rejected Buber's radical view that there are no constraints on the kind of thing towards which one could form an I-Thou relation. (He famously discusses the example of an I-Thou relation with a tree, for instance (Buber 1958, pp. 7-8).) The motivation for rejecting that line of thought was to preserve the intuitive appeal of Buber's description of dialogical relations as relations between people. The agapic doctrine just specified now helps to see the basis of that appeal. Accordingly, love is a defensible response to another person's personhood, even though that personhood alone is not sufficient to justify loving them.

As I have said, my goal is not to defend the rationality of agape, so I will not argue that the doctrine just outlined is correct. Rather, my view is that this doctrine captures the form that the agapic relation must take, whether or not it ultimately belongs in an ethical outlook. In order to defend this view, I will discuss the deficiency of either episode without at least the promise of the other. I will then also contend that Buber himself had in mind a similar hybrid solution to the problem.

Suppose first that agape could consist solely of the stance of openness, without ever involving moments in which the loving party focused on the characteristics of the other nor ever appraised them to be, in themselves, worthy of the love. This is the Kantian view, versions of which I have referred to already as the bare personhood view, as espoused by the likes of Velleman and Setiya. It rejects the agapic doctrine as stated above in favour of the idea that one has sufficient reason to love any particular person, simply in virtue of their being a person. As such, love is construed as a perennial openness to the other's individual characteristics. The problem with agape consisting of this episode without the other is that it is rendered hollow, in a sense. That is, in merely *projecting* warmth from the subject's side of the relation, the agapic-lover would refuse to engage with the distinctive features of the other, would refuse to acknowledge that about them in particular which makes them special: their qualities, charms, achievements, and so on. Refusing to acknowledge such things, whilst acting

---

[81] Thanks to Shelly Kagan for suggesting this formulation.

with beneficence towards someone, is quite far from a loving attitude. Indeed, it is a strikingly cold-hearted way to treat someone – never to deign to engage with what makes them who they are, but to keep the other's individuality always at arm's length.

To illustrate this, consider again the faltering collegial relationship between Arwen and Dilys. We left Arwen in a predicament: she faced a request from Dilys that she could see no reason to accede to, at least not for as long as she looked for such a reason in Dilys' rather nasty-seeming character. I suggested that *Openness* offered a solution to this predicament because according to that conception of agape, the way to act lovingly towards Dilys is precisely not to look too closely at her character, but to bestow love upon her regardless, to direct the love at Dilys *herself*, shorn of her apparently unpleasant traits.

If Arwen were to act with agapic openness towards Dilys on this occasion, by agreeing to swap their days off, we would hardly be inclined to say on the basis of this one kindly deed alone that Arwen has *loved* Dilys. And the hollowness of *Openness* becomes apparent when one imagines this deed as part of an extended series of such acts of indiscriminate kindness. That is, imagine that every time Arwen interacted with Dilys she acted selflessly in Dilys' interest, helping Dilys in every way within her means, but all the while remaining absolutely uninterested in Dilys' character. Such disinterest, it seems, is incompatible with truly loving another. If Dilys' attempts to improve her conduct in the office, to act more kindly towards Arwen, and less snobbishly in general, were met with no recognition whatsoever from Arwen, then far from feeling loved, Dilys may well feel that her relationship with Arwen has become almost completely estranged.

As a criticism of the Kantian, or bare personhood views, it might seem that my point here is overly simplistic. I am suggesting that they fail to account for the nature of love because in being indiscriminate, the conception of love those views advance is cold-hearted, and thus omits the personal warm-heartedness that is central to distinctly loving attitudes. This suggestion may indeed seem simplistic since proponents of the Kantian view are not naïve to the fact that their universal conception of love seems cold, and impersonal, in comparison to rival conceptions, such as the qualities view. Kantians advance their impersonal conception of love, in spite of its apparent coldness, because only such a view which conceives love as directed towards personhood as such, they think, can explain the universality of love as a moral attitude, and one which it is appropriate to hold towards people in general. However, the episodic conception of agape that I am proposing, as expressed in the doctrine stated above, does explain exactly that universal appropriateness of agape. By incorporating the rational role of personhood without accepting a conception of love as perennial openness, the episodic theory of agape shows that the cold-heartedness of the Kantian view is not necessary.

So if agape were to consist in a stance of openness alone, then it would be hollow, and would not really be love at all. But what if agape were to be understood simply as the stance of attention, and never of openness? To some degree, this possibility has already been addressed with Arwen's initial predicament. Attending to the love-worthy traits of another is a necessary component part of what it is

to agapically love them. But that attitude cannot be sufficient for agape for the simple reason that agape must move one to act lovingly towards even those with whose love-worthy traits one is not yet adequately acquainted. As Murdoch's example of the mother and daughter-in-law illustrated, paying attention to another can take time. But if, until one has learnt to see the underlying merits of another's character, one treats them with corresponding disdain (or whatever other neutral or negative evaluative attitude is presiding), one cannot be said to love that other. When Arwen was first asked to swap days off with Dilys, she could see nothing in Dilys that warranted kindness – her attention to Dilys' character was at that stage, at best, embryonic. So if agape were merely a matter of devoting attention to others' characters and acting on the basis of one's subsequent appraisals, then Arwen would have refused Dilys' request.

As such, agape as a love of people in general must be comprised of both episodes: first of openness, and then continuing with attention. Interestingly, the idea of love as something that oscillates between episodes of openness, and episodes of appraising the beloved, is something that Buber already mentions explicitly in *I and Thou*. Putting a little more detail on the episodes that Buber himself had in mind in this regard, will, I hope, represent a helpful development of the model of agape that I am proposing.

As mentioned above, the first and fundamental idea in *I and Thou* is that there are two stances that a subject can take towards things: a stance with a determinate object, an object with conceptually articulable characteristics; and a stance which doesn't have such an object. The former Buber calls the '*I-It* primary word', the latter the '*I-Thou* primary word'. As above, I do not mean here to give a very faithfully Buberian account of agape, so for my purposes, the *I-Thou* stance towards someone can be understood to be a non-judgemental attitude of openness. Conversely, the *I-It* stance towards someone can be understood just as one of appraisal and evaluation. Relevantly to the model of agape that I am proposing, Buber thinks of love as an attitude that is generated when the subject relates to the other as *Thou*, but which can nonetheless be sustained when the subject's judging faculties are reintroduced.

Expressing that genesis of interpersonal love in the stance of the *I-Thou* primary word, Buber says of the beloved:

> [W]ith no neighbour, and whole in himself, he is *Thou* and fills the heavens. This does not mean that nothing exists except himself. But all else lives in *his* light. (1958, p. 8, emphasis in original)

This aphoristic claim appeals to the thought that in its essence, love has as its object the person themselves, irrespective of any transient, contingent qualities. It is a similar thought to that which motivates figures like Kreft (under review), Setiya (2014), and Velleman (1999), to suppose that the object of all love is always in this sense non-qualitative, that love is never based on a judgement of the beloved's individual characteristics. But Buber disagrees. To be sure, he thinks it is possible – perhaps only momentarily – to relate to a person in this pure fashion, where one's stance towards them is not clouded by some qualitative description of who and what they are. But these moments are necessarily ephemeral:

> [T]his is the exalted melancholy of our fate, that every *Thou* in our world must become an *It*. It does not matter how exclusively present the *Thou* was in the direct relation. As soon as the relation has been worked out or permeated with a means, the *Thou* becomes an object among objects – perhaps the chief, but still one of them, fixed in its size and its limits. (Buber, 1958, p. 16-17)

And given this exalted melancholy – that we cannot persist in our *I-Thou* relations with others – what becomes of our love of those others? To this question, Buber is unequivocal:

> [L]ove itself cannot persist in direct relation. It endures but in interchange of actual and potential being. (1958, p. 17)

The love that is formed in relation to a person as a *Thou*, endures when that person becomes an *It* for the loving subject. That is, the valuing-attitude that was formed by embracing the person as a whole – in a mood of their 'filling the heavens' – that attitude lives on even when the other ceases to fill the heavens quite as they had done. It lives on even when the judging, evaluating attention of the subject is turned upon the beloved, with its helplessly critical eye. Thus, this attitude that endures the transition from openness to critical attention must be one that can be seen, from the latter point of view, as justified: as an apt response to the evaluable features of the object. When the other who was regarded as a *Thou* becomes an *It*, the subject who loved them as *Thou* must look at their individuating characteristics in search of signs of a worthiness of the love that has been bestowed. In this way the subject can maintain an attitude of loving the other across these polarised stances that she takes towards the beloved.

Conclusion

I have argued that it is not obvious what kind of relation agapic love consists in. This is because while, on the one hand, it might seem to be right to relate to people with a kind of non-judgemental openness, on the other hand, people seem to deserve to be attended to by others, and to have their valuable qualities seen and appreciated. I have suggested that these two ways of relating to people can both figure as complimentary episodes together constituting the structure of agapic love. Agape is a *pair* of stances. And from an episodic combination of these stances, one can always be moved to undertake supererogatory altruistic deeds in the interests of the other. When one is acquainted with their charms, virtues, and whatever other love-worthy qualities they may bear – sufficiently acquainted to be motivated to do them good, out of appreciation of those qualities – then one's stance towards that other is, justly, one of simply attending to their character and developing that acquaintance with those qualities. Contrastingly, when one interacts with others without any such acquaintance – as with strangers and one's enemies – people one is inclined to dislike for example – treating them as a person requires something more. It requires abandoning one's judgement of them, relating to the person themselves, as a whole person, and loving them for who they are. This kind of love, though, in order to survive the subject's transition from the stance of openness to the stance of judging, must be rational in form. It is a love that *calls for* a justification by reference to the love-worthiness of the individual characteristics of the other, albeit a justification that must come retroactively.

In proposing a model of what agape *is*, I have omitted to give full consideration to how it is normatively grounded. That remains a question for another time. As such, the issue of whether agape can be made complete sense of from a secular perspective has not been even putatively concluded here. Indeed, justifying the notion of interpersonal openness that I have drawn from Buber looks like a philosophical puzzle. Despite this, I hope to have displayed some of the appeal of the idea of *I-Thou* relations as a model of how it is sometimes possible for people to relate to one another. Thereby, I also hope to have evinced the relevance of Buber's work to these issues in contemporary philosophy.

Since Openness draws so directly from the notion of a dialogical stance – that is, the stance of addressors and addressees relating as such to one another – it is clear that agapic love on the Openness model has a second-personal character. Furthermore, in Chapter Three, I showed how the features of others that could justify loving them might be epistemically accessible only within second-person relations. On that basis, the Attention side of the discussion in the present chapter can also be thought of as having a certain second-personal character. As such, the present chapter has not been a dialectic between a second-personal and a *third*-personal conception of agape. Rather, the dialectic takes place between two different ways in which such a loving stance might depend upon the second-person relation. Indeed, it is my suggestion that this kind of moral, universal love has a structure that is shaped, episodically, in both of these second-personal ways.

Chapter 6

Second-personal encounters and value epistemology

In this chapter I will take up a rather different approach to trying to explain the practical significance of the second-person relation. Previous chapters have discussed the distinctive kinds of reason-giving that might be possible only when agents relate to one another second-personally, and, similarly the distinctive kinds of reasons that arise from such relations. These discussions have represented a piecemeal approach to appreciating the significance of the second-person relation. That is, I have tried to show, through analysis, how certain particular kinds of reasons that figure in our lives could depend on distinctly second-personal relations between people. This piecemeal approach was partly motivated on the back of argument in Chapter Two that there is reason to doubt Darwall's grander, wholesale approach, which attempts to show, in one fell swoop so to speak, how *all* moral obligations are dependent on second-person relations. However, despite my criticisms of Darwall's own wholesale approach, this chapter returns to consider an alternative argument that would find second-person relations to be pervasively significant for all moral reasons. Rather than this significance being, *à la* Darwall, a conceptual connection between second-personal attitudes and moral obligation, the focus here will be on the *epistemic* role that such relations might play in morality. I will develop a line of thought that was introduced in my schematic discussion in Chapter Three, by considering the idea that second-person relations provide epistemic access to values which form the very basis of agents' practical outlooks: values which thus play a role in grounding a large array of altruistic reasons. In this way, second-personal relations could be crucial in gaining access to reasons that might otherwise be thought to be, in a sense, third-personal.

Specifically, the thought that will be under investigation – which is drawn from the work of Emmanuel Levinas – is that a second-personal encounter can occasion a kind of cognition of the intrinsic value of the life and ends of another person. In the first section below I will motivate a problem in value epistemology. Then in the second, I will set out some characteristics of Levinas' account of what he calls *the face*. Eschewing certain other readings of this account, I will pursue an interpretation which says that the ethical significance of the face of the other is epistemological. The face does not (or does not only) command, or request, or demand, or summon; it teaches. A general challenge for such an interpretation is what to make of the idea that we could learn something from the 'otherness' of the other, which is, *ex hypothesi*, beyond our grasp. In the third section below, therefore, I will go into greater detail explaining the role that the confrontation with the other's otherness can play in the value-epistemological theory on offer. My primary goal is not so much to offer a conclusive defence of the thesis under discussion, but merely to explain its plausibility and

illustrate some advantages that it has over nearby rival views in moral epistemology. An auxiliary goal is to go some way in defending the attribution of this view about value-epistemology to Levinas.

1. Locating the intellectual context

Within a particular strand of meta-ethics is a problem to which I think the early Levinas has an interesting answer. The question is this: how can people have knowledge of intrinsic value outside of themselves? Levinas' answer, in short, is that such knowledge can be gained through reflection, but a sort of reflection that is prompted by experience. In order to animate the force of this question, a number of points must be clarified – the first of which is the concept of intrinsic value, as distinct from extrinsic value.

Something is extrinsically valuable to the extent that it is valuable in virtue of its relation to other things. A person's life might be thought to be extrinsically valuable to a certain extent in virtue of the role they and their life play in valuable social practices, where the practices, with their social roles, confer a certain kind of value on the lives of those who occupy them. Perhaps, for example, your sister has a certain sort of value for you just in virtue of being your sister. This is not, however, the most familiar sense in which a person's life is thought to be valuable. We commonly think that the most important sense in which a life is valuable is the sense of its being *intrinsically* valuable. This, by contrast with extrinsic value, just means the value a thing has independently of its relation to other things. Acknowledging the intrinsic value of another person's life means acknowledging the value their life has *in itself*.[82]

There is no special epistemic worry about knowledge of extrinsic values. Judging something to be extrinsically valuable involves deploying an existing normative conviction, and assessing its relation to natural, which is to say non-normative, facts. Judgements of extrinsic value always *assume* some normative premise, add some further non-normative premise(s), and thus draw a normative conclusion. Suppose, for example, that one values rice pudding in virtue of its relation to one's appetite, namely, its propensity to bring about such pleasing experiences when eaten. How does one gain epistemic access to the fact that rice pudding is extrinsically valuable? Well, one starts by assuming that pleasurable eating experiences are valuable (either intrinsically or extrinsically), and

---

[82] It is by now widely thought in value theory that the intrinsic / extrinsic distinction should be kept quite separate from the non-instrumental / instrumental distinction. See (Korsgaard 1983) for a defence of this separation. In any case, in this chapter I am focusing on the value of persons, so the non-/ instrumental distinction cannot be directly relevant. This is because the kind of thing that can be valued instrumentally is a state of affairs, namely, a state of affairs is instrumentally valuable if it brings about another state of affairs that has (either instrumental or non-instrumental) value. But the value of a life, and indeed the value of the liver of a life, are not states of affairs. As such, they are not the kinds of things that could be valued in the instrumental way. When A is treated by B as a mere means to B's ends, A is not valued as herself by B at all, not even instrumentally. In such an instance, it is not A that B values as a means, but some action of A's, or some event that A can help to bring about. This is all to emphasise the point that for reasons above and beyond value theorists' qualms about a conflation between extrinsic and instrumental value, for present purposes it is clearly only the former that is of interest.

then one learns, through experience, that rice pudding possesses properties of being delicious. Such a properties are natural facts, and whatever epistemological difficulty there may be in explaining how one can come to know them is not any more pertinent regarding the formation of beliefs about extrinsic value than it is regarding the formation of any other kind of belief about the world.

There is, however, a special kind of epistemic worry regarding knowledge of intrinsic value: namely, how can we ever know that something is intrinsically valuable? Whereas facts about extrinsic values can be expressed in terms of descriptions of the natural world (with further *assumed* facts about what is valuable), intrinsic values are irreducibly normative. In order for one's belief that X is intrinsically valuable to count as knowledge, it must be, at least, true and justified. A central issue of meta-ethics concerns the matter of whether such beliefs can be true. It is in response to something like this question that non-cognitivists deny that intrinsic valuing is a matter of believing. Nihilists and error-theorists respond in turn by accepting that valuing is indeed based in beliefs, but denying the existence of the values that those beliefs are putatively about. In the face of these challenges from non-cognitivists, nihilists and error theorists, this issue seems a rather general one, and I am not proposing that Levinas can be read as providing a novel answer to the question on such a general construal.

There is, though, another level at which the question can be understood. That is within a *realist* approach to meta-ethics, which maintains both that valuing is at least partly an attitude of holding beliefs about valuable things, and that such valuable things do exist. Within this approach, the question of how we can ever know that something is intrinsically valuable is not about whether such beliefs could ever be true (which is bracketed for the time being as a metaphysical rather than epistemological issue), but rather about whether they could ever be *justified*.

In this context, the challenge is to explain how there can be facts about what is valuable that bear a relation to human minds, where that relation is such that we can grasp those facts. The most prominent answer to this challenge is known as *intuitionism* and was pioneered by the British Moralists. Intuitionism is the view that people can rationally intuit, or apprehend, irreducibly normative facts, such as intrinsic values. Intuitionists claim that some propositions are such that when we as rational beings merely contemplate those propositions, we can know them to be true without any further reasoning, because this set of propositions are self-evident.[83]

At this point it is important to note a significant difference between the kind of ethical intuitionism that I am going to elaborate upon and attribute to Levinas, and the kinds of intuitionism traditionally advanced by the some of the proponents of intuitionism in the British Moralist tradition including Richard Price (1974) and A.C. Ewing (1970). Their version of intuitionism focuses on moral principles as the kind of propositions that they take to be self-evident. There is scope for variation over which principle, or principles, they think to be self-evident and thus knowable in this

---

[83] For a helpful summary of ethical intuitionism, see Philip Stratton-Lake's (2002) Introduction to his edited collection on the matter.

distinctive way. A committed consequentialist, for example, could think that the principle that one can rationally intuit is a consequentialist principle that a right action is the one which brings about the best consequences. What distinguishes the Levinasian view that I will develop, from this tradition, then, is its focus not on principles, but on propositions about values.

As will become clear shortly, focusing on propositions about value rather than about principle will help the Levinasian version of intuitionism to address an objection. For now, note that on the face of it propositions about value are just as important as propositions about principles in grounding an ethical outlook. Consider whatever ethical principle one might think to be fundamental and knowable by being self-evident – the principle of consequentialism, for example. In order to be applicable to the world in such a way that it produces moral judgements, the consequentialist principle requires any agent who is applying it to believe that others – and not only themselves – are morally valuable beings whose lives and experiences ought to be taken into consideration as part of the goodness of the outcomes of any given action. So whilst intuitionism has historically focused on a method for coming to know about principles, it might just as well have focused on propositions about values.

In any case, intuitionism is committed to epistemological foundationalism. That is to say, the apprehension of an irreducibly normative fact must be non-inferential. If it were an inference, it would have to be based on a premise of the form 'X is [say] intrinsically valuable if X has property Y'. But this premise itself already assumes a normative fact, namely that property Y instantiates intrinsic value. And how could one come to know *this* fact, by inference, without presupposing yet another normative fact as an antecedent? There is thus a vicious regress threatening the view that normative facts could be apprehended inferentially, so rational intuitionism posits a kind of non-inferential apprehension, which provides an epistemic foundation to normative thought.

Intuitionists thus invite us to suppose that it is possible to grasp intrinsic values intuitively. Think, for instance, of one's own life. If asked whether one believes that one's own life is valuable, not extrinsically but in itself, psychologically healthy people will, I imagine, answer affirmatively. And if one is then asked *why* one believes this, on what basis, in virtue of what one takes one's own life to be intrinsically valuable, one might perhaps be able to offer reasons in support of this conviction, but such reasons will bottom out quite soon. Let us suppose that in fact one has no further reasons and that the proposition at which one's justifications bottoms out is the very proposition that one's own life is intrinsically valuable. Any individual entertaining this belief can reasonably believe it without needing a justification. With regard to each of our own lives, from own individual perspective, we *just do* have this belief.[84] Which is to say, precisely, that we hold the belief on the basis of intuition, on the basis of the proposition's self-evidence.

---

[84] Not everyone believes their own life to be intrinsically valuable. Indeed, many people with depression do not. My point here does not depend on the universality of this belief, but rather on the immediacy and foundational quality that the belief typically has for those that do hold it.

2. Motivating a problem

For realists, it is important that the reasons that cause us to have a normative belief coincide with the reasons that justify those beliefs. If there is a normative fact that my life is intrinsically valuable, but my belief in the same proposition was caused erroneously, say by the indoctrination of the narcissistic culture into which I was raised, or if it was hard-wired in me by evolution (Street 2006), then my belief is only true by luck. That belief thus does not amount to knowledge, and is not a sound basis for my normative outlook. But for each of us this belief we have in the value of our own lives is so central to our outlook and all of our emotional responses that it is difficult to doubt that the causal reason for our holding the belief is not any undue influence of culture or of the views of those around us. The cause of our belief is the very intuition which, according to intuitionism, can provide us with knowledge of the normative fact in question.

Things are a little less straightforward when it comes to values external to any given agent's own life. Here, a worry could be entertained that what causes one to believe that others' lives are valuable things is the mere custom of one's community, a view one was brought up into, taught by one's parents, and so on. The belief that other's lives are valuable seems in some ways less immediate and less indubitable than the value of one's own life. So if it was the fact that belief in the value of one's own was based in intuition that gave it that immediacy, then the absence of the same immediacy in the case of others' lives prompts a worry that the basis of the belief may be different.[85]

To be sure, it is not a problem for realism as such if *sometimes* one's axiological beliefs in values outside of oneself are caused erroneously, that is, by factors other than the cognitive apprehension of those values by the subject. What would be a problem is if such beliefs were *always* caused erroneously. As such, what intuitionism requires is an explanation of how such cognition can sometimes take place, so that it can in principle vindicate our value-outlooks.

This brings us back to the question of how we can have knowledge of intrinsic value *outside* of ourselves. And this is the question to which I propose that Levinas' account of the face of the other can be read as an answer. As I will now go on to argue, the encounter with the face of the other is best understood as being ethically significant by being epistemologically significant. In the face, one comes into indubitable contact with a value outside of one's own life, which one apprehends non-inferentially, and which thus provides an epistemic foundation for a moral axiological outlook.

---

[85] I hope that my expression in this paragraph of the difference between the significance for an agent of their own life, and of the lives of others, avoids both the 'objectionable individualism', and the 'objectionable collectivism' that Lavin (2014, p. 282) identifies in this domain. That is, on the one hand, the kind of egoist approach to practical philosophy which regards the value of others' lives as utterly mysterious and removed from any given agent's concerns, is by now widely discredited. On the other hand, that discredited individualism contained at least a kernel of truth, which is what I have tried to articulate above: that there is *some* difference between an agent's sense of their own value, versus the value of another. Incidentally, a further thought in (Lavin 2014) is that both Darwall and Korsgaard have too general a notion of the second-person to be able to capture the significance of second-personality for ethics. The Levinasian viewpoint expressed below is offered as a corrective to exactly this short-coming, although perhaps in a different spirit to Lavin's own perspective in that paper.

3. The phenomenology of the face

I have motivated a problem in value epistemology. I propose now to offer a reading of Levinas' account of a morally significant interpersonal encounter that he calls the face, which addresses that problem. The face is a term that recurs a great deal in Levinas' earlier writing, where it is a central motif in *Totality and Infinity* (Levinas 1969), as well as other writings written in the years either side of that work. Here is a passage which invokes a cluster of his core ideas regarding the encounter.

> Faith is not the knowledge of a truth open to doubt or capable of being certain; it is something outside of these modalities, it is the face to face encounter with a hard substantial interlocutor, who is the origin of himself, already dominating the forces which constitute him and sway him, a you, arising inevitably, solid and noumenal, behind the man known in that bit of absolutely decent skin which is the face, which closes over the nocturnal chaos and opens upon what it can take up and for which it can answer. (Levinas 1987, p. 41)

In unpacking this passage, there are three main ideas to be delineated which together comprise Levinas' theory of the face. In an encounter with the face of the other, the subject:

i)   relates to the other second-personally;

ii)   experiences the 'otherness' of the other;

        and,

iii)   is presented with a reason to act morally.

I will explain the motivation behind each of these ideas presently. But first, some further clarity is needed over the scope of Levinas' claims, not just in this particular passage but in his ethical phenomenology generally.

There is some ambiguity about the status of this description. One can construct a spectrum of readings of this description, ordered in a sequence from being, at one end, strongly committed to certain psychological and metaphysical claims, to, at the other end, weaker and relatively non-committal readings. Thus, it is not completely clear whether: (a) Levinas' description of the encounter is intended to capture what it is *always* like to experience embodied interpersonal interaction; or (b) only to capture what it is always like to experience such interaction *when* there is a mutually addressed, second-personal stance taken up by both parties; or (c) whether the account is only of a one-off formative moment that each of us must have at some point to become a morally mature agents; or (d) whether (at the weakest end of the spectrum) this is just a description of how interpersonal encounters can sometimes be, but not necessarily how they ever in fact are.

The first construal, (a), is deeply implausible. There are many interpersonal interactions in which the other's otherness or their ethical significance do not occur to the subject, even momentarily. Examples of this include, for instance, acts of blatant inconsiderateness: people are capable, sometimes, of interacting with others without being moved at all by the interests of those others. So the phenomenology should not best be understood as a necessary description of all interpersonal life.

The second option, (b), is also unattractive because it faces a Scylla and Charybdis of circularity and implausibility. On this construal, Levinas' phenomenology is understood as saying that

all second-personal interaction involves an experience of encountering the other's otherness and, concomitantly, being presented with a reason to act morally. This could be understood as a definition of what it is to relate to someone second-personally. That is all very well, but since it is a definition it cannot be all that informative so is a rather disappointing way of making sense of what is putatively the flagship idea of Levinas' early philosophy. On the other hand, this construal could itself be construed as a synthetic claim, that all second-personal relations (defined in some other terms) involve this ethical encounter with the otherness of the other. But as a matter of fact, this seems unlikely, for similar reasons to those that dismissed (a). If second-personality is understood broadly as the kind of relation within which people find it appropriate to make addresses to one another using second-personal pronouns, then there are surely many such relations which are not experienced as ethical encounters, and where neither party contemplates the other's otherness. Examples will include exchanges in which people use second-personal pronouns while perpetrating ugly violence against their addressees.

The third option, (c), is a tempting one in the context of the over-arching structure of *Totality and Infinity*. One way of reading the trajectory of that work is as a kind of *ontogenetic* story: a general characterisation of how an individual subject makes the transition from self-centred egoism to mature moral agency, via an encounter with the face of the other. I do not wish to completely dismiss this reading, and perhaps the ontogenetical perspective is one that Levinas had in mind at various points.[86] But it is not the construal that I would like to recommend here, and for what it is worth, it does also seem unlikely. In this case, the unlikelihood of Levinas' phenomenology as an ontogenetic account stems from its methodology. The matter of how people in fact come to form and exercise moral concepts, incorporating due consideration of others' interests into their practical reasoning, is an empirical matter. It is one that will ultimately be determined by developmental psychology, not armchair philosophy.

So I have put aside options (a), (b), and (c) on the spectrum of readings of Levinas' description of interpersonal experience. That leaves option (d): that this is just a description of how interpersonal encounters can sometimes be. To many Levinas enthusiasts, this may seem somewhat underwhelming. If all the account of the face does is articulate a kind of experience that is merely possible, one that only *might* occur, then that account is not obviously of huge philosophical significance. But this is nonetheless the correct reading in my view. And even though this is at the modest end of the range of ways that such moral phenomenologies might be understood, it does nonetheless represent a bold and significant putative contribution to moral philosophy. The account of the face does this by providing a description of how moral beliefs – of the intrinsic value of others' lives – could in principle be well-founded.

---

[86] For discussions of that ontogenetical perspective in Levinas, see (Bahler 2015) and (Vandenberg 1999).

With this clarification settled, let us return to the quoted passage and the aforementioned triad of core ideas to be found there. The essence of the theory of the face is, once again, an experience in which a subject relates second-personally to an other, encounters their otherness, and is thereby presented with a reason to act morally. In calling this the essence of the theory, I omit consideration of a number of other important themes that Levinas often draws together with the face, including notions of height, vulnerability, nudity, and transcendence. But because of the interdependence of Levinas' core terminology during this phase of his work, I think that those terms too could be drawn in to this same interpretation of what he thinks is going on when a subject encounters the face of another.

One important aspect of what he thinks is going on is that the encounter with the face causes a *rupture* in the structure of the subject's ordinary experience. Whereas the subject is accustomed to treating everything outside of itself as a mere thing – a matrix of objects to be conquered, negotiated, enjoyed – the face of the other is different. This structural difference is that between the customary third-personal mode of experience, and the second-personal mode. Thus in the passage quoted, the other is referred to as an 'interlocutor,' and 'a you'. Regarding the other as so totally different from the rest of the world outside of the subject is a crucial component of the epiphany that goes along with recognising the face. Rather than being a passive object with which the subject could do as they please, the other is an interlocutor. By relating to the other as a second-person, the subject is committed – in the very structure (or *mode*) of their experience – to treating as practically significant the *intentions*, and *will* of a being other than themselves.

So the third-personal mode of the subject's ordinary experience is ruptured. It is interrupted by the presence of something – someone – recognised from the second-person perspective of an interlocutor. And it is in this context that the subject becomes aware of something outside of its grasp – an otherness beyond comprehension. That is to say, the otherness of the other is something of which the subject becomes aware only when they relate to the other second-personally. It is illuminating to notice why this is. Why, for instance, could the subject not have grasped the other's otherness *third*-personally? Why must the encounter with the face be second-personal?

The answer is that Levinas understands the second-person relation to be partly comprised of a stance of openness towards the other. Following Buber, Levinas describes the stance of a 'you', or of an interlocutor as one in which the mysterious depth of the other becomes salient. That hiddenness of the other's interior world becomes salient because as an interlocutor, the other is a being whose responses can never be taken for granted. Indeed, the other's every contribution to their second-personal interaction with the subject is such that it could never be exhaustively understood in advance, by the subject. Unlike the objects that populate the world in the third-personal mode of experience, which can (and must) be understandable and predictable in order to be conquered and enjoyed, the second-person is, to some extent, a law unto themselves. If one truly treats another as an interlocutor, then one is necessarily open to them uttering things that one could not expect. So the mind of the other to which one relates second-personally – as the source of such unexpected possibilities – lies

necessarily outside of the subject's complete comprehension. That is the sense in which this otherness of the other could not be encountered if the other were treated third-personally.

That otherness is what Levinas refers to variously in the quoted passage as the other being the 'origin of himself', which is to say that the other's will (confrontationally) does not originate from the subject's own. The other is alternatively described as being 'noumenal', which is to say that the content of the other's mind and their interior world is in principle beyond the scope of the phenomenal realm, it cannot be exhaustively grasped. Similarly, Levinas makes reference in the passage to that which is 'behind the man', and there a 'nocturnal chaos'. Again, this vocabulary emphasises the fact that when the subject encounters the face of the other, they are confronted with the awareness that the other contains a great deal that lies beyond the subject's understanding.

In the next section below I will offer some elaborations of some different ways in which one might understand Levinas' thought about the ungraspable other presented in the interpersonal encounter. For now, what is relevant is that the ungraspability does confront the subject, and it is in connection with this that the subject is gripped by a moral reason to act in the interest of the other. That is, just as the inexplicability of the other follows for Levinas from the second-personality of the face-to-face interaction, so the ethical dimension follows from that inexplicability.

A big interpretative question at this point, then, is *why* this recognition, that the other is so other that some great part of them is ungraspable to the subject, should give rise to any reason for the subject to act in the other's interest. In what remains of this section I will warn against a too-easy answer to that interpretative question, and offer a first pass at what I think is a better answer. In the more detailed discussion of the nature of the other's otherness in the next section, I will elaborate on this better answer.

A tempting line of interpretation is one that connects up Levinas' thought about the moral significance of the face with notions of empathy and sympathy. But this would be a mistake. Consider Levinas' mention in the quoted passage to the 'absolute decency' that the other's face embodies. This expression recalls a range of other expressions that Levinas uses in *Totality and Infinity* to stress the ethical dimension of the face. One might think, then, that this ethical dimension is *just* the *feeling* that we do often indeed have when we look into the face of another, of wanting to help them. That is, it is known that recognising another person's face can trigger sets of psychological phenomena such as empathy, sympathy, pity and compassion. To feel pity, or compassion, or sympathy, is to see the face of another and to be moved to help them, or at least to feel as though one ought to help them. And since this experience seems to match up with what Levinas is describing in the encounter with the face, it seems to add intuitive attraction to his account that it matches is up with the phenomena of empathy and sympathy.

However, there are several reasons to reject the connection between empathy or sympathy, and the experience that Levinas articulates in the account of the face. For one thing, the encounter with the face stirs a kind of *awe* in the subject. The realisation that the other is so totally external to

114

the subject, and different from the subject, serves to establish the other's dignity and respectability, and indeed, their 'height' over the subject. These sentiments are in tension with the condescending sentiments of sympathy, pity and compassion.[87] In the latter, the subject always regards themselves as being simply in a position of power over the other, whereas in the kind of interpersonal experience that Levinas has in mind, the power balance is more complicated than that. The other is *both* below the subject – in need – and at the same time above the subject – standing over the subject as a judge (Stern 2018, p. 13).[88]

Moreover, a further and even more important reason to dissociate the face from the experience of sympathy is that the ethical significance of the face must be understood as having some actual *de jure* impact on the subject. That is, it must be *right* for the encounter with the face of the other to make a difference to the subject's practical reasoning. Empathy and sympathy, insofar as they are reduced to psychological mechanisms, cannot alter the subject's normative situation. If it is true that when one looks into another person's eyes one is likely to feel more inclined to act beneficially towards them, then this fact is about a (possibly hard-wired) feature of human psychology, not about some normative reasons that depend on whether or not one is looking into the other's eyes. Thus, construing the encounter with the face as an experience of empathy or sympathy threatens to eliminate the normative component of Levinas' story and equate with a normatively-neutral descriptive account of what it in fact *feels like* one should do when one looks into the face of another.

For these reasons, the encounter with the face should not be linked too closely with experiences of empathy or sympathy. This suggestion serves as a general reminder of the difficulty of making sense of Levinas' account as being both descriptive and, in some way, normative. For if he intends to be describing merely what it is like to interact in a certain way with another person, then how could it follow from this that one ought to act a certain way? I have already indicated how I think this question should be answered, namely as follows: the normative impact of the encounter with the face of the other is that it opens up the possibility for the subject to *know* that the other's life is of value.

To show that this is a theoretically significant answer to the value-epistemological problem raised in section one, it is worth spelling out how the beliefs prompted by the encounter might be thought to count as knowledge. In short, the belief in the intrinsic value of the other that arises in the encounter counts as knowledge because as well as being true, it is formed exactly on the basis of an appreciation of the independence of the value of the other, which means that it can be justified and

---

[87] Despite this *prima facie* tension, Levinas ultimately wants to explain the ethical force of the face of the other in a way that reconciles the 'dimension of height', with the 'dimension of depth'. For an interpretative attempt to show how such a reconciliation might go, see (Lewis & Stern, forthcoming).

[88] To be clear, my concern in this chapter pertains to the epistemology of the 'dimension of depth', not the 'dimension of height'. That is, my focus is on how we come to know that other's lives are intrinsically valuable, not on how we come to know that others have authority over us. It should be possible to investigate these epistemological issues independently from one another even though, as Levinas emphasised, the two matters seem to be closely related.

reliably caused. As such, Levinas' account of the face offers a story of value epistemology that assuages two pervasive worries around the notion of coming to know of intrinsic values.

The first worry that Levinas' account avoids was alluded to earlier: that the belief in the other's intrinsic value could be formed as a mere product of social inculcation – some kind of ideological conviction passed down through the influence of others. But the encounter with the face is an epiphany. It is not just a rupture with the ordinary mode of experience, but with the ordinary framework of the subject's beliefs. That is to say, for the subject to find themselves gripped by the value of a being totally outside of itself comes as a surprise. This kind of surprise is exactly what is needed to explain the possibility of actual knowledge against the worry that the belief in the other's intrinsic value could have been socially inculcated. Joel Kupperman (2005, p. 677) notes a similar point, that moments of surprise about the existence or extent of some intrinsic value are crucial to any plausible explanation of how one's beliefs about their value could escape the worry of social influence.

The other worry that Levinas defuses concerns the object of the cause of the belief. I mentioned in section one that in order for any of us to think that our beliefs are well-founded, we have to have certain beliefs about our beliefs. In particular, we have to believe that our beliefs are caused in some appropriate way by that which makes them true. Thus, if one is a rational intuitionist about moral epistemology, then one must believe that one's belief in some foundational moral fact is in some way caused by that moral fact. One can see that as compared with the Levinasian approach just articulated, this is a more pressing problem for traditional approaches to rational intuitionism. Their focus, traditionally, has been on absolute moral principles as exemplary of the kind of moral fact that one might rationally intuit, which would then provide a sound epistemic basis for one's other moral beliefs: principles such as the categorical imperative, or the principle of utility. The problem of course is in explaining how the fact of the matter about these principles could be the cause of one's belief in them.

Levinas shifts the focus of rational intuitionism. His version of this moral-realist epistemology is directed not at general moral principles, in the first instance, but at the particular intrinsic value embodied by a particular other.[89] In this context, it is easier to see how the relevant

---

[89] It should be noted here that there is some *prima facie* textual resistance to interpreting Levinas as claiming that the encounter with the face involves the encounter with the *value* of the other: he says on *TI*, p. 202, 'The presentation of being in the face does not have the status of a value'. But it is quite clear in the context of the passage that what Levinas means here by 'the status of a value' is something like the status of a socially constructed fact, a truism to be passed around from one person to the next, but which could turn out to be nothing more than a societal prejudice. That kind of socially constructed fact resides, Levinas says, in 'the ambiguity of the true and the false which every truth risks – an ambiguity, moreover, in which all values move' (*TI,* p. 202). And certainly for Levinas, the ethical significance of the other that the subject encounters in the face is nothing of this sort. It turns out, therefore, that Levinas' claim here sits quite comfortably with my reading, for while the other's ethical significance is not a mere value in the sense that Levinas has in mind, it is of course a kind of value in another sense. That is, this significance which the subject directly encounters, in a way that cannot be immediately doubted, is such that it warrants valuing-attitudes of concern, and is such that the subject is motivated and warranted to act in its interest. On a more neutral conception of the term than the one Levinas uses at this point in the text, this is just what it is for something to be a value.

beliefs could be appropriately caused. Namely, the belief that this particular other's life is of intrinsic value is caused by an appreciation of the sheer independence of their subjective world from that of the subject. This appreciation triggers (in a way that I will discuss in the next section) the rational intuition that their life is valuable, intrinsically, which is the proper basis for such a belief. So in being an epiphany, and in being directed at an appropriate object, the beliefs formed in the encounter with the face escape the two pertinent challenges to the justification of axiological beliefs that otherwise plague rational intuitionism.

4. Epistemology and otherness

I have now sketched an interpretation of Levinas' ethical phenomenology as offering an answer to a challenge in a particular strand of value-epistemology. Before concluding, it will now be fruitful to return to elaborate on a central concept within Levinas' account, with a view to articulating more detail in this epistemological picture: namely, the concept of 'otherness'. What exactly is it? And can more be said about what it has to do with the other's intrinsic value? I will draw out and distinguish between three ideas for explaining what the other's 'otherness' consists in, and how it might be epistemologically significant.

A general remark to make about the concept of otherness, is that for Levinas it means more than merely being different from the subject. By which I mean, the other person cannot be understood simply in terms of their differences from the subject. Indeed, the concept of otherness does not denote 'other than the subject' so much as 'other than that which the subject can incorporate into a framework of understanding' (or a *totality*). Despite this, Levinas plays on the ambiguity of the words 'other' and 'otherness' which do not in themselves determine what they are other *to*. In light of this, as noted, I think there are three senses of otherness at play in Levinas' value epistemology:

i)   the other's individual characteristics;
ii)   the other's own practical perspective;
        and,
iii)   the other's intrinsic value.

I will explain what is meant by each of these, and argue that in Levinas' epistemological story, recognition of (i) and (ii) are together the epistemic enabling conditions at play in the subject's realisation of (iii), in the encounter with the face.

First, then, consider the other's individual characteristics. These are of course 'other' in the sense of marking the other apart from the subject. So at first blush they are not the sort of thing that Levinas is referring to with his notion of otherness. Indeed, to the extent that the other's individual traits can be rendered conceptually intelligible to the subject – as some such traits certainly can be – they can be incorporated into a totality of understanding from the subject's point of view. And moreover, there is a sense in which the subject can try to overcome the other's otherness by reducing

the other to a complex definition of specified individual characteristics. Thus, once someone has been exhaustively described, from the dimensions of their physical appearance to every aspect of their history and their discernible psychological profile, they will no longer really be so terribly 'other' at all.

Importantly, though, not all of an other's individual characteristics are immediately clear to a subject who is confronted with this other in an embodied interpersonal encounter. Rather, many of the judgements that the subject cannot help making about the other are almost completely psychologically opaque to the judging subject. In each perceptual mode there are a great many details to be taken in that the subject immediately judges, but does so sub-consciously in a way that they cannot make explicit immediately – and which perhaps they *never* can.

For example, in the visual field, when encountering an other, the subject is prone to make opaque judgements about a range of things: the other's gait, and the manner of their other bodily movements; the history that is disclosed by the lines and general comportment of the other's face, as well as their facial micro-expressions, the direction of their gaze, and so on. Every twitch of a muscle and momentary lingering of a finger in the air forms a datum in the subject's apprehension of the other. But these details are far too complex for the subject to interpret consciously. So the interpretation instead takes place subconsciously, and as a result is only partly accessible to conscious reflection.

What is more is that the visual field seems to be the perceptual mode that provides the *most* consciously articulable kind of experience. The other is also perceived and appraised through other modes. Their voice is always judged by the subject, sub-consciously, in its myriad details: aspects of the other's accent, the emotional underpinnings of their intonation, their timbre, the varying melodic lines of their speech, and the size and depth of its pauses.[90] Indeed, the other is also immediately experienced through olfactory and tactile modes – which give rise to experience that are famously difficult to bring to the surface of articulable conceptual thought.

Relevantly, these traits of the other's that the subject subconsciously judges through the variety of perceptual modes, are not merely traits that individuate the other in a purely descriptive sense. Many of them also play a role in forming an evaluation of the other. People manifest their more-or-less love-worthy characters through many of these subconsciously appraised expressions. In fact, when our judgement of someone being worthy of love is difficult to articulate, this inarticulability plausibly owes something to the opaque quality of the subconscious judgements that form the greater portion of our awareness of other people.

So insofar as this myriad of judgements are not within the subject's consciously articulated framework of understanding, they are 'other' in Levinas' sense. A great bulk of the characteristics which individuate the other are received as an emotionally (because evaluatively) charged *feeling* of

---

[90] The subtleties of people's appraisals of one another's humanity through their voices is a subject that is explored in (Cavarero 2005).

otherness. That is one part of what the subject encounters when they counter the face. But it plays a role in making it possible for the subject to also appreciate further aspects of the other's otherness.

The next such aspect worth considering is the other's own practical perspective. By appreciating the other as a being who could be worthy of such interpersonal valuing-attitudes as love, and as one who is individuated in such a complex bundle of characteristics as to render them unique, the subject is pressed to regard the other as an agent in their own right. This in turn enables the paradigm shift discussed above from third- to second-personal orientation. Only by thinking of the other as an agent with their own subjective perspective can the subject treat the other as an interlocutor or a 'you'.

For the other to have their own practical perspective amounts to several things. It means that they are a locus of affective and sensory experience – that is, they are sentient. So part of what the subject recognises in the face is the other's sentience. And the sentience that the subject attributes to the other has a significant dimension of otherness to it, because whilst the subject assumes that the other's experience of joy and suffering, and of sights and sounds, are in some way analogous to the like experiences of the subject's own, this is coupled with the inescapable acknowledgement that the other's experiences, being theirs, will be different. And the manner in which they are different is in principle epistemically inaccessible. Part of what the subject recognises in the face of the other is that there is something that it is like to be the other, and that the subject can never know what this something is.

That recognition of an individual locus of sentience is crucial, and it is connected to the further recognition of the other's agency – because the other's doxastic and axiological perspective are their own in a sense that is analogous with how their sentient experience is their own. And the subject realises this in the encounter with the face. Just as the subject has one set of beliefs about things, including about what matters and how much, so the other has their own. Just as the subject has a certain set of capacities that fix what they can *do*, so the other has their own. And just as what it is like to feel and see things from the other's perspective is epistemically inaccessible to the subject, so too is their agential perspective. The subject can never know exactly how strongly the other really holds any given belief of normative conviction, nor how exactly various values are weighed up against each other within the other's outlook. Neither can the subject ever gain more than a superficial insight into the dynamic flux within which the other's beliefs and intentions and values travel through stages of commitment and rejection from the other's point of view.

I have tried to put some flesh on the bones of the first and second senses – (i) and (ii) – of how the otherness of the other is encountered in the face. I think that acknowledging the other's otherness in these senses is what Levinas thinks sufficient for the subject to rationally intuit the third sense (iii): the other's intrinsic value. This is not to say that the subject forms an inference of the form, 'if X is other in senses (i) and (ii), then X is intrinsically valuable'. This is because, as noted earlier in this chapter, rational intuitionism must posit some *non*-inferential belief at the foundation of

normative epistemology. So (i) and (ii) are not premises in an inference to (iii). They are rather just the enabling conditions within which (iii) can be apprehended.

This may seem like a rather *ad hoc* connection between the components of Levinas' view of the face, but in fact it is deeply rooted in his outlook. For Levinas, the experience of the face is a rupture in the ordinary structure of experience. The nebulous awareness of that which differentiates the other, and the concomitant realisation of the inaccessibility of the other's practical perspective, are the active parties which break the subject out of their ego-centric, third-personal, subject-object mode. But Levinas makes perfectly clear in multiple places that the encounter with the face is not *just* a negative break with the ordinary sort of experience, it is also the positive beginning of another. It is in this spirit that the passage quoted in section two speaks of the 'faith' that is identical with the face-to-face encounter. The face is an epiphany that makes possible 'Language' and 'Justice' (*TI,* pp. 78-79, 100-1), for Levinas. He regards the encounter as the moment that opens up the possibility for intersubjective community where moral concepts can be exchanged and rational cohabitation can be communally deliberated. The domain of rational social life is premised, for him, on the sharing of basic normative propositions: namely, the undeniable value of the life of each, not as instruments to some greater end, but in and of themselves.

Conclusion

In this chapter I have offered an interpretation of Levinas' account of the encounter with the face of another. I have argued that the best way to read this account is as a moment in which the subject is presented with an *epistemic* reason to act morally, namely, on the basis of the intrinsic value of the other person. When conceived this way, Levinas' account can be seen to make a worthwhile contribution to meta-ethical debate about moral cognition. His picture provides several resources to enrich and defend a rational intuitionism about irreducibly normative facts. In short, those resources are: the shifting of focus from the cognition of absolute moral principles to the cognition of particular intrinsic values; the rich phenomenological account of how intrinsic values, rationally intuited, can surprise us; and the explanation of how our beliefs about such values can plausibly be brought about in the normatively appropriate way, which is by apprehension of their proper object.

What this discussion further provides is, in certain respects, a theory of how the second-person relation is significant in moral reasoning *tout court*. Just as Darwall, channelling Fichte, found moral normativity to rely on the second-person perspective, so does Levinas. But whereas Darwall claims that it is in virtue of the possibility of second-personal demands that moral reasons have obligatory force, Levinas' alternative shows second-personal relations to be epistemic enabling conditions for well-founded moral convictions.

Conclusion

In concluding this dissertation, I would like to draw to the surface and make explicit some common threads that run through it, which define the theoretical approach that I have been advancing. To do so, I will offer a summary of what I have done, narrating the main ideas that have occupied the six preceding chapters. But first, it will be instructive to mention something that I have *not* discussed in great detail, a perspective in opposition to which the spirit running through my project can best be understood. Namely, I have not directly taken up the Fichtean line of thinking which sees the second-person perspective to be part of the structure of first-personal thought and agency.

In *Foundations of Natural Right*, Fichte says (2000, p. 37) that 'if there are to be human beings at all, there must be more than one'. This is the conclusion of an argument which he takes to establish that social life is a necessary condition for individual subjectivity. Indeed, the kind of social life that is required for such freedom is distinctly second-personal in form: it is 'free, reciprocal interaction' with others which is constitutive of our individual freedom, according to Fichte. On this line of thought, then, the second-person perspective is a formal, or structural, part of what it is to think free thoughts, or undertake free actions as an individual. That is, when *I* form a belief or an intention, my very doing so is implicitly connected to the possibility of reciprocal interaction with a *you*. My beliefs and intentions are free, and they are mine, because I am *answerable* for them, where the justifications that imbue those attitudes with their status as free can only make sense as justifications in the context of being questioned, as though from the perspective of a second person.

This summary of that line of thinking is not a very faithful reconstruction of how Fichte argues in *Foundations of Natural Right* – which anyway is not my aim – but it is a fairly natural way of arriving at the same, Fichtean point: that the first-person perspective itself depends on the second-person perspective. Now, in acknowledging that this dissertation has omitted any extensive consideration of this line of thought, I should also note that it *has* been mentioned in places, if only as a peripheral focus. For, Fichte's thought about the second-person relation has been of much interest in recent philosophy, including parts of recent philosophy that I have discussed. In my survey of conceptions of the second-person in Chapter One, I mentioned a number of contemporary writers who are strongly influenced by Fichte, including Sebastian Rödl (2014), Adrian Haddock (2014), and, not least, Darwall himself (2006, esp. chpt. 10). Beyond merely noting these views in my introductory survey, Chapter Two was devoted to exploring and criticising Darwall's views, and by extension, his version of the Fichtean approach to understanding second-personality. So it is not as though the Fichtean approach has been completely overlooked in the forgoing.

Nonetheless, one might think it incumbent on an enquiry into the practical significance of the second-person relation to devote its central focus to this increasingly popular perspective which sees the second-person as having practical significance *everywhere*, in all free thought and action. What I

would like to point out now, however, is that it is no accident that my central focus has been elsewhere. The Fichtean line of thinking regards the significance of the second-person relation to be unitary – a form that pervades practical thought and exerts the same kind of structural influence everywhere it exerts any influence at all. In direct contrast to that view, a core feature of my project – especially in Chapter Three – has been to bring to light the *diversity* of ways in which second-person relations can legitimately influence action. One might indeed suspect that a focus on whatever unitary, structural significance the second-person relation might have could actively obscure that diversity of other significances.

Moreover, because the Fichtean role for the second-person relation in practical reasoning is a structural role – forming part of the nature of individual reasoning – it is impervious to the differences between the influence that can be borne by agents, qua second persons, in different kinds of relationships. Because it sees no difference between the second-person standpoint occupied by an enemy, or that occupied by a friend, the Fichtean approach is, in this regard, cold-hearted and detached. Again, my project has been opposed to exactly such a cold-hearted approach to understanding how others can guide our action when we relate to them as you. A driving thought in several of the later chapters of this work has been that second-person relations open up ways for warm-hearted interpersonal sentiments – interpersonal valuing-attitudes in my terms – to manifest in reasons for action. As I argued in Chapters Three and Four, one important site at which second-person relations come together with interpersonal valuing-attitudes is the set of normative powers that such valuing-attitudes generate to be exercised second-personally. These are exemplified by the power to make requests. Requesting is a power which subjects grant to certain others to create discretionary reasons for those subjects to act, where the reasons created are partly grounded in the discretionary value of the requester for the requestee.

So my work has been developed in contrast to the Fichtean approach to the same enquiry in two senses: in emphasising the diverse range of ways that the second-person relation can have a bearing on action, and in appreciating the role that second-personal relations can yield practical significance *through* warm-hearted interpersonal relationships. The latter contrast is connected to a third, and this third contrast is perhaps the most salient point at which my approach has been defined against what I am labelling as the Fichtean approach: namely, the *generality* of the second-person standpoint as a structure of individual agency on the latter approach, which contrasts with the *particularity* of second persons on my approach, in the practical influence they can exert through second-person relations.

This contrast is an echo of the tension that I presented in Chapter Two that resides within Darwall's discussion of the second-person standpoint. On the one hand, one can think of the second-person relation as between a subject and a faceless, nameless other whose questions and demands – or the mere possibility thereof – make the subject into the accountable, responsible, free agent that she is. That is the Fichtean thought that Darwall develops. On the other hand, the view that I have been

recommending attends to the significance that particular others have when (and only when) we relate to them second-personally. In Chapters Three, Five and Six, I advanced a raft of arguments to the effect that when we are entwined in second-person relations with others, then, and then alone, can we come into contact with that about them in particular that warrants certain practical attitudes from us. These attitudes might be loving, or caring, or valuing in some other moral or non-moral way. What is important from my perspective, but overlooked by Fichte's, is how our practical lives are shaped by actual, particular others with whom we relate as an *I* to a *you*.

Bibliography

Abramson, K. & Leite, A., 2011. Love as a Reactive Emotion. *The Philosophical Quarterly*, 61(245), pp.673–699.

Adorno, T., 2005. *Minima Moralia: Reflections on a Damaged Life*, London: Verso.

Andolsen, B.H., 1981. Agape in Feminist Ethics. *Journal of Religious Ethics*, 9(1), pp.69–83.

Anon, 2011. *The Holy Bible: 400th Anniversary Edition of the King James Version*, New York: Harper Collins.

Bahler, B., 2015. Levinas and the Parent–Child Relation: A Merleau-Pontyian Critique of Applying Levinas's Thought to Developmental Psychology. *The Humanistic Psychologist*, 43(2), pp.128–147.

Blum-Kulka, S., House, J. & Kasper, G., 1989. *Cross-Cultural Pragmatics : Requests and Apologies*, Ablex Publishing Corporation.

Bohman, J., 2000. The Importance of the Second Person: Interpretation, Practical Knowledge, and Normative Attitudes. In H.-H. Kögler & K. R. Stueber, eds. *Empathy and Agency: The Problem of Understanding in the Human Sciences*. Westview Press, pp. 222–242.

Brandom, R., 2000. Facts, Norms, and Normative Facts: A Reply to Habermas. *European Journal of Philosophy*, 8(3), pp.356–374.

Brandom, R., 1994. *Making It Explicit : Reasoning, Representing, and Discursive Commitment*, Harvard University Press.

Brandom, R., 2007. The Structure of Desire and Recognition. *Philosophy & Social Criticism*, 33(1), pp.127–150.

Brinck, I., 2008. The Role of Intersubjectivity in the Development of Intentional Communication. In J. Zlatev et al., eds. *The Shared Mind: Perspectives on Intersubjectivity*. Amsterdam: John Benhamins Publishing Company, pp. 115–140.

Buber, M., 1958. *I and Thou* 2nd ed., Edinburgh: T. & T. Clark.

Cavarero, A., 2005. *For More Than One Voice : Toward a Philosophy of Vocal Expression*, Stanford University Press.

Conant, J. & Rödl, S., 2014a. Introduction. *Philosophical Topics*, 42(1), pp.1–12.

Conant, J. & Rödl, S., 2014b. The Second Person. *Philosophical Topics*, 42(1).

Cupit, G., 1994. How Requests (And Promises) Create Obligations. *The Philosophical Quarterly*, 44(177), p.439.

Dancy, J., 2004. *Ethics Without Principles*, Oxford University Press.

Darwall, S., 2011. Authority, Accountability, and Preemption. *Jurisprudence*, 2(1), pp.103–119.

Darwall, S., 2012. Bipolar Obligation. In R. Shafer-Landau, ed. *Oxford Studies in Metaethics Volume 7*. Oxford University Press, pp. 333–358.

Darwall, S., 2013a. *Honor, History, and Relationship : Essays in Second-Personal Ethics II*, Oxford: Oxford University Press.

Darwall, S., 2016. Love's Second-Personal Character: Reciprocal Holding, Beholding and Upholding. In E. E. Kroeker & K. Schaubroeck, eds. *Love, Reason, and Morality*. Abingdon: Routledge, pp. 119–135.

Darwall, S., 2013b. *Morality, Authority, and Law : Essays in Second-Personal Ethics I*, Oxford University Press.

Darwall, S., 2007. Reply to Korsgaard, Wallace, and Watson. *Ethics*, 118(1), pp.52–69.

Darwall, S., 2004. Respect and the Second-Person Standpoint. *Proceedings and Addresses of the American Philosophical Association*, 78(2), p.43.

Darwall, S., 2006. *The Second-Person Standpoint : Morality, Respect, and Accountability*, Cambridge, MA: Harvard University Press.

Darwall, S., 2017. Trust as a Second-Personal Attitude (of the Heart). In P. Faulkner & T. Simpson, eds. *The Philosophy of Trust*. Oxford University Press, pp. 35–50.

Darwall, S., 1977. Two Kinds of Respect. *Ethics*, 88(1), pp.36–49.

Davidson, D., 1992. The Second Person. *Midwest Studies in Philosophy*, 17(1), pp.255–267.

Driver, J., 2014. Love and Duty. *Philosophic Exchange*, 44(1).

Economidou-Kogetsidis, M., 2010. Cross-Cultural and Situational Variation in Requesting Behaviour: Perceptions of Social Situations and Strategic Usage of Request Patterns. *Journal of Pragmatics*, 42(8), pp.2262–2281.

Eilan, N., 2014. The You Turn. *Philosophical Explorations*, 17(3), pp.265–278.

Enoch, D., 2014. Authority and Reason-Giving. *Philosophy and Phenomenological Research*, 89(2), pp.296–332.

Enoch, D., 2011. Giving Practical Reasons. *Philosophers' Imprint*, 11.

Ewing, A.C., 1970. *Reason and intuition.*, London: Haskell House Pub Ltd.

Faulkner, P., 2014. The Moral Obligations of Trust. *Philosophical Explorations*, 17(3), pp.332–345.

Feuerbach, L., 1972. *Principles of the Philosophy of the Future*, www.marxists.org. Available at: https://www.marxists.org/reference/archive/feuerbach/works/future/index.htm.

Fichte, J.G., 2000. Foundations of Natural Right. *Cambridge texts in the history of philosophy*, p.xxxv, 338.

Frankfurt, H.G., 2004. *The Reasons of Love*, Princeton University Press.

Gallagher, S. & Hutto, D.D., 2008. Understanding Others Through Primary Interaction and Narrative Practice. In J. Zlatev et al., eds. *The Shared Mind: Perspectives on intersubjectivity*. Amsterdam: John Benhamins Publishing Company, pp. 17–38.

Grau, C., 2004. Irreplaceability and Unique Value. *Philosophical Topics*, 32(1), pp.111–129.

Habermas, J., 2000. From Kant to Hegel: On Robert Brandom's Pragmatic Philosophy of Language. *European Journal of Philosophy*, 8(3), pp.322–355.

Haddock, A., 2014. On Address. *Philosophical Topics*, 42(1), pp.345–350.

Heal, J., 2014. Second Person Thought. *Philosophical Explorations*, 17(3), pp.317–331.

Heck, R.G.J., 2000. Do Demonstratives Have Senses? *Philosopher's Imprint*, 2(2).

Hieronymi, P., 2008. The Reasons of Trust. *Australasian Journal of Philosophy*, 86(2), pp.213–236.

Huangfu, W., 2012. Cross-Cultural Comparisons of English Request Speech Acts in Native Speakers of English and Chinese. *Cross-Cultural Communication*, 8(4), pp.24–29.

Jeske, D., 2008. *Rationality and Moral Theory: How Intimacy Generates Reasons*, London: Routledge.

Jollimore, T., 2011. *Love's Vision*, Princeton: Princeton University Press.

Jollimore, T., 2015. The Importance of Whom We Care About. In A. Rudd & J. Davenport, eds. *Love, Reason, and Will: Kierkegaard After Frankfurt*. London: Bloomsbury, pp. 47–72.

Jones, K., 1996. Trust as an Affective Attitude. *Ethics*, 107(1), pp.4–25.

Kádár, D.Z. & Haugh, M., 2013. *Understanding Politeness*, Cambridge: Cambridge University Press.

Kant, I., 1997. *Groundwork of the Metaphysics of Morals*, Cambridge: Cambridge University Press.

Kierkegaard, S., 2013. Kierkegaard's Writings, XVI: Works of Love. , p.584.

Kolodny, N., 2003. Love as Valuing a Relationship. *The Philosophical Review*, 112(2), pp.135–189.

Korsgaard, C.M., 2007. Autonomy and the Second Person Within: A Commentary on Stephen Darwall's The Second-Person Standpoint. *Ethics*, 118(1), pp.8–23.

Korsgaard, C.M., 1983. Two Distinctions in Goodness. *The Philosophical Review*, 92(2), p.169.

Kraut, R., 1987. Love De Re. *Midwest Studies in Philosophy*, 10(1), pp.413–430.

Kreft, N., Love and the Moral Aspect of Irreplaceability.

Kukla, R. & Lance, M.N., 2009. *"Yo!" and "Lo!" : The Pragmatic Topography of the Space of Reasons*, Harvard University Press.

Kupperman, J.J., 2005. The Epistemology of Non-Instrumental Value. *Philosophy and Phenomenological Research*, 70(3), pp.659–680.

Laing, R.D., 1960. *The Divided Self (An Existential Study in Sanity and Madness)*, London: Penguin.

Lance, M. & Kukla, R., 2013. Leave the Gun; Take the Cannoli! The Pragmatic Topography of Second-Person Calls. *Ethics*, 123(3), pp.456–478.

Lavin, D., 2014. Other wills: the second-person in ethics. *Philosophical Explorations*, 17(3), pp.279–288.

Lavin, D., 2008. Review of Stephen Darwall, The Second Person Standpoint: Morality, Respect, and Accountability. *Notre Dame Philosophical Reviews*, 1. Available at: https://ndpr.nd.edu/news/the-second-person-standpoint-morality-respect-and-accountability/.

Levinas, E., 1987. *Collected Philosophical Papers*, Dordrecht: Martinus Nijhoff Publishers.

Levinas, E., 1969. *Totality and Infinity: An Essay on Exteriority*, Pittsburgh: Duquesne University Press.

Lewis, J.H.P., 2018. The Discretionary Normativity of Requests. *Philosopher's Imprint*, 18(20), pp.1–16.

Lewis, J.H.P. & Stern, R., Commanding, Giving, Vulnerable: What is the Normative Standing of the Other in Levinas? In M. Fagenblat, ed. *Levinas and Analytic Philosophy: Metaethics, Personal Identity, and the Moral Life*. London: Routledge.

Liao, S.M., 2007. The Idea of a Duty to Love. *The Journal of Value Inquiry*, 40(1), pp.1–22.

List, C. & Pettit, P., 2011. *Group Agency*, Oxford University Press.

Ludwig, K., 2016. *From Individual to Plural Agency*, Oxford University Press.

Markovits, J., 2014. *Moral Reason*, Oxford University Press.

Martin, M.G.F., 2014. In the Eye of Another: Comments on Christopher Peacocke's "Interpersonal Self-Consciousness." *Philosophical Studies*, 170(1), pp.25–38.

McMyler, B., 2011. *Testimony, Trust, and Authority*, Oxford University Press.

Mill, J.S., 1998. *Utilitarianism* R. Crisp, ed., Oxford University Press.

Montaigne, M. de, 1993. *Essays*, London: Penguin.

Murdoch, I., 1970. *The Sovereignty of Good*, London: Routledge.

Nygren, A., 1982. *Agape and eros*, University of Chicago Press.

Outka, G., 1977. *Agape: An Ethical Analysis*, New Haven: Yale University Press.

Owens, D., 2006. A Simple Theory of Promising. *Philosophical Review*, 115(1), pp.51–77.

Owens, D., 2012. *Shaping the Normative Landscape*, Oxford University Press.

Peacocke, C., 2014. Interpersonal Self-Consciousness. *Philosophical Studies*, 170(1), pp.1–24.

Postema, G.J., 2001. Law as Command: The Model of Command in Modern Jurisprudence. *Philosophical Issues*, 11(1), pp.470–501.

Price, R., 1974. *A Review of the Principal Questions in Morals* D. D. Raphael, ed., Oxford: Clarendon Press.

Protasi, S., 2016. Loving People for Who They Are (Even When They Don't Love You Back). *European Journal of Philosophy*, 24(1), pp.214–234.

Raz, J., 1999. *Practical Reason and Norms*, Oxford: Oxford University Press.

Raz, J., 1988. *The Morality of Freedom*, Oxford University Press.

Rödl, S., 2014. Intentional Transaction. *Philosophical Explorations*, 17(3), pp.304–316.

Rödl, S., 2007. *Self-Consciousness*, Harvard University Press.

Rorty, A.O., 1987. The Historicity of Psychological Attitudes: Love Is Not Love Which Alters Not When It Alteration Finds. *Midwest Studies in Philosophy*, 10(1), pp.399–412.

Satne, G., 2017. Brandom and the Second Person. *International Journal of Philosophical Studies*, 25(2), pp.189–209.

Satne, G., 2014. What Binds Us Together. *Philosophical Topics*, 42(1), pp.43–61.

Scheffler, S., 2012. *Equality and tradition : questions of moral value in moral and political theory*, Oxford University Press.

Schilbach, L. et al., 2013. Toward a Second-Person Neuroscience. *Behavioral and Brain Sciences*, 36(4), pp.393–414.

Schroeder, M., 2006. Cudworth and Normative Explanations. *Journal of Ethics and Social Philosophy*, 1(3).

Searle, J., 1975. Idirect Speech Acts. In P. Cole & J. L. Morgan, eds. *Syntax and Semantics, Vol. 3: Speech Act*. New York: Academic Press, pp. 59–82.

Searle, J.R., 1969. *Speech acts : an essay in the philosophy of language*, Cambridge University Press.

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

Sheinman, H., 2011. *Promises and Agreements*, Oxford University Press.

Stern, R., Levinas, Darwall, and Løgstrup on Second-Personal Ethics. In M. L. Morgan, ed. Oxford University Press.

Stern, R., 2019. *The Radical Demand in Løgstrup's Ethics*, Oxford: Oxford University Press.

Stratton-Lake, P., 2002. *Ethical intuitionism : re-evaluations*, Oxford University Press.

Strawson, P.F., 1968. Freedom and Resentment. In *Studies in the Philosophy of Thought and Action*. London: Oxford University Press.

Street, S., 2006. A Darwinian Dilemma for Realist Theories of Value. *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition*, 127, pp.109–166.

Thompson, M., 2004. What is it to Wrong Someone? A puzzle About Justice. In R. J. Wallace et al., eds. *Reason and Value: Themes from the Moral Philosophy of Joseph Raz*. Oxford: Clarendon Press, pp. 333–384.

Vandenberg, B., 1999. Levinas and the Ethical Context of Human Development. *Human Development*, 42(1), pp.31–44.

Velleman, J.D., 1999. Love as a Moral Emotion. *Ethics*, 109(2), pp.338–374.

Wallace, R.J., 2012. I—R. Jay Wallace: Duties of Love. *Aristotelian Society Supplementary Volume*, 86(1), pp.175–198.

Wallace, R.J., 2007. Reasons, Relations, and Commands: Reflections on Darwall. *Ethics*, 118(1), pp.24–36.

Wallace, R.J., 2019. *The Moral Nexus*, Princeton: Princeton University Press.

Wanderer, J., 2010. Brandom's Challenges. In B. Weiss & J. Wanderer, eds. *Reading Brandom: On Making It Explicit*. London: Routledge, pp. 96–114.

Weil, S., 1952. *Gravity and Grace*, London: Routledge and Kegan Paul.

Weil, S., 2009. *Waiting for God* J. M. Perrin, ed., New York: Harper & Row.

Williams, B., 1981. *Moral Luck: Philosophical Papers, 1973-1980*, Cambridge University Press.

Zahavi, D., 2001. Beyond Empathy: Phenomenological Approaches to Intersubjectivity. *Journal of Consciousness Studies*, 8(5–7), pp.151–167.

Zahavi, D., 2014. *Self and Other: Exploring Subjectivity, Empathy, and Shame*, Oxford: Oxford University Press.