


Love First  
by  
Patrick Quinn White



Submitted to the Department of Linguistics and  
Philosophy  
in partial fulfillment of the requirements for the degree of  
Doctor of Philosophy  
at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2019

© Massachusetts Institute of Technology 2019. All rights reserved.

Author .. **Signature redacted**

Department of Linguistics and Philosophy  
August 28, 2019

Certified by .. **Signature redacted**

Kieran Setiya  
Professor  
Thesis Supervisor

Accepted by ... **Signature redacted**  
/ .....  
Bradford Skow  
Chair, Committee on Graduate Students



# Love First

by

Patrick Quinn White

Submitted to the Department of Linguistics and Philosophy  
on August 28, 2019, in partial fulfillment of the  
requirements for the degree of  
Doctor of Philosophy

## Abstract

How should we respond to the humanity of others? My dissertation argues that the fundamental answer is love, defending and carrying out the beginnings of a love-first approach to ethics.

The ideal of love for all (agape) has religious origins, and in my first chapter, I show that it can be detached from this religious context and still serve as a foundation for ethics. Ordinary love we have for friends and family is subject to “outward pressure”: in our love for the few, we find reason to love all. I explain how we can make sense of such love, even love for those we have not met, in terms of the neglected phenomenon of “plural love,” as when we love the members of our family, as such. The problem with agape is not that it is impossible but that it conflicts with the good of selective love. In light of this conflict, I argue that we should treat agape as an ideal to be approximated. We can understand respect as the minimum required approximation of love and derive the basic features of deontological ethics—commitment to equality, respect for autonomy, non-aggregative concern for well-being, imperfect duties of beneficence—from the ideal of agape.

Where the first chapter is an investigation into how we should act towards others in general, the second turns to partiality, beginning with our love for friends and family. Suppose I help my friend Kevin, acting out of love, when I might have spent my time and resources helping strangers. It seems obvious that I am justified in being partial because we are friends. Yet when I act out of love, I am not motivated by the fact that we are friends (which seems like “one thought too many”) but by facts like Kevin’s need. The normative and motivational desiderata on a theory of partiality thus pull in opposite directions: where the relationship seems to play a justifying role, it

need not be a motivating reason. The solution lies in recognizing the neglected diachronic nature of partiality, and of practical rationality. To act of out love is not to be moved by the fact of one's relationship; it is to have a history of so relating to someone. And it is in virtue of that history that non-relational facts (e.g., that Kevin needs help) have greater rational significance. History counts not as a reason on which we act but as a phenomenon that affects the weight of other reasons. I derive this fact from a (defeasible) rational requirement of constancy and show how this requirement can be generalized to other cases of practical reasoning in the face of parity.

My third chapter asks when and why we should tell the truth. Philosophers have traditionally answered this question with an exclusive focus on lying and deception, ways of obscuring the truth, while giving little thought to indiscretion—sharing or eliciting a truth that should be left unsaid. I argue that these vices and their correlative virtues must be theorized together. A unified account of honesty and discretion must start with the concrete relationships between speaker and interlocutor. Our relationships determine what information is and is not “in bounds.” These communicative norms are constitutive of our relationships—that my friend can ask about my private life where my colleague cannot is part of what makes our relationship an intimate friendship. I argue that our reasons to tell the truth are explained by our relationships: we have reason to follow their communicative norms just insofar as we have reason to be in them. We moreover have reasons to share or withhold the truth in order to shape our relationships: we can put certain topics in or out of bounds, molding our relationships into something new.

Thesis Supervisor: Kieran Setiya  
Title: Professor

## Acknowledgments

To thank all of those who have helped me in completing this dissertation seems nearly impossible. I am no doubt omitting many who have made an enormous difference. But most of all, I would like to acknowledge those whose encouragement, criticism and patient conversations are responsible for the best of what is here: Kieran Setiya, Caspar Hare, Jack Spencer, Tamar Schapiro, Daniel Muñoz, Brad Skow, Simon Keller, Brendan Dill, Kevin Dorst, Agustín Rayo, Nathaniel Baron-Schmitt, Ryan Preston-Roedder, Dimitri Halikias, Ela Nagele, Tyler Wilson, audiences at MIT's MATTI and thesis workshops and those at NNEPA. And of course my family and friends without whom I would not have made it: Brian White, Carmen Medina, Yoseph Desta, Alex Trimm, and my partner, Maily Fidler. Thank you all.



# Contents

1	Ethics in the Shadow of Love	9
1.1	Introduction . . . . .	9
1.2	A quasi-account of love: feeling and doing . . . . .	14
1.3	Love's Outward Pressure . . . . .	18
1.4	Plural Love . . . . .	26
1.4.1	Particular love and plural love . . . . .	27
1.4.2	Love for the Spencers . . . . .	29
1.5	Agape . . . . .	35
1.5.1	Universal and Impartial Regard . . . . .	37
1.5.2	Wellbeing and Autonomy . . . . .	41
1.5.3	Anti-Aggregation . . . . .	48
1.5.4	Summing Up . . . . .	49
1.6	Agape as an ideal . . . . .	49
1.6.1	Agape dominates a life . . . . .	50
1.6.2	Ideals . . . . .	54
1.6.3	Respect as an Approximation of Agape . . . . .	56
1.7	Love First . . . . .	59
2	The Importance of Being Constant	65
2.1	Partiality . . . . .	65
2.2	Taking Love Seriously . . . . .	69

2.3	Relationships Views . . . . .	70
2.3.1	Relationships views disallow directly motivated action . . . . .	72
2.4	Individuals Views . . . . .	75
2.4.1	Getting the Moral Psychology Right . . . . .	76
2.4.2	Getting the Ethics Wrong . . . . .	78
2.5	An Appeal to History . . . . .	85
2.6	Constancy as a General Requirement . . . . .	90
2.7	Details and Objections . . . . .	94
2.7.1	Depth and History . . . . .	94
2.7.2	Getting stuck and Bootstrapping . . . . .	97
2.8	Conclusion . . . . .	100
3	Honesty and Discretion . . . . .	103
3.1	Shifting the Approach . . . . .	106
3.1.1	The Explanandum: Motivating Integrationism . . . . .	107
3.1.2	Method: Motivating Localism . . . . .	113
3.2	A Relationships View . . . . .	120
3.2.1	Information and Communicative Practices . . . . .	122
3.2.2	Relationships & Constitutive Ends . . . . .	126
3.3	Which Relationships Matter . . . . .	132
3.3.1	Flawed Relationships . . . . .	137
3.4	Why Be Truthful? Why Discreet? . . . . .	144
3.5	A Coda on Lying . . . . .	146



# Chapter 1

## Ethics in the Shadow of Love

“We need a moral philosophy in which the concept of love, so rarely mentioned now by philosophers, can once again be made central.”

---

Iris Murdoch, *The Sovereignty of Good*

### 1.1 Introduction

How should we respond to the humanity of others? Ethics is in large part about that—we inhabit a world filled not only with mountains and rivers and streams, but with people, and it can be hard to know what to do about them. In secular Western ethics the responses have largely fallen into one of two camps: first, that people have interests and we should promote those interests, e.g., by making as many of them as happy as possible; and second, that people are autonomous beings whose capacity for free choice must be respected. It is sur-

prising that what might be the most popular answer to this question dropped out of the running. Buddhism, Judaism and Christianity, whose adherents together amount to a plurality of human beings in history, exhort us to love others. The commandment to love one's neighbor as oneself is at heart of the Decalogue.<sup>1</sup> Paul goes further, saying "The entire Law is fulfilled in a single decree: 'Love your neighbor as yourself'. (Galatians 5:14)"<sup>2</sup> This ideal of agape is taken to form the basis not only of relations among friends, family and fellow Christians, but to all, even one's enemies (Matthew 5:35-46). The Buddhist ideal of Bodhicitta is of love for all sentient life. As the Dalai Lama puts it, "the liberation of mind by love is practiced with universal pervasion by extending it to all beings, then all breathing things, all creatures, all persons, and all those with a personality" (Dalai Lama XIV and Chodron, 2014).

To be clear, it is not as though philosophers, historical or contemporary, have ignored love. The charge that a given ethical view fails to make room for love is a powerful one that was given perhaps most famous voice by Williams (1981).<sup>3</sup> And recently, the literature on love has grown tremendously, tackling questions about what love is, what our reasons for love are, and how we can fit loving, partial relationships in to our broader ethical pictures. But unlike Christians and Buddhists for whom an ideal of love for all serves as the central, organizing ethical notion, most of these philosophers see love as confined to the domain of intimate relationships between friends, family,

---

<sup>1</sup>See Goodman (2008) for an excellent discussion of the love commandment in the Jewish Tradition.

<sup>2</sup>Or Matthew 22:38-40, where Jesus says all the law of the old testament can be summed up in only two commandments, to love God and to love one's neighbor as oneself.

<sup>3</sup>See also Stocker (1976).

romantic partners and the like. “Outside the context of close relationships, love can be ignored. It is not at the root of obligation, as such, but an aspect of our personal lives that calls for moral reflection. On [the dominant] approach [to love], one could have an adequate view of the nature and grounds of morality without having much to say about love” (Setiya, 2014).<sup>4</sup>

And, you might think, with good reason! Love has its home in our special relationships with the few. Ethics is concerned with much more than that, and it is hard to see how some kind of ideal of love could guide our actions toward the many, especially if we divorce the notion from the religiously inspired metaphysics that prop it up. First, there are a host of normative questions. Why think of love for all as an ideal? Granted that we might have some reason to respect everyone or to care about their interests; what reason do we have to love everyone? It is striking that virtually all of the reasons are religious: God’s love for us is agapic,<sup>5</sup> and so we should imitate him; Jesus commanded us to; it is only through love of all that one can achieve union with all and the dissolution of the self.<sup>6</sup> And for that matter, might not love of all be a bad thing? Sure, the thought that ethics is about love sounds nice, but is it really? As Freud wrote, “A love that does not discriminate seems to me to forfeit a part of its own value” (2010, p. 82), and Wolf’s famous description of “the loving saint” helps bring to life that one who really lives her whole life out of a kind of boundless love for others seems not to be living an ideal life (1982). Love for all would dominate a life; with so many people in so much need, it is hard to see how one who loves all could do anything

---

<sup>4</sup>To be clear, Setiya does not endorse the dominant approach.

<sup>5</sup>Some use “agapeistic” (outka), others “agapistic.” With apologies for doing some violence to the Greek but in the interest of fewer syllables, I’ll use “agapic.”

<sup>6</sup>Outka and others.

but serve others—indeed, it is hard to see how she could maintain a space for intimate, special relationships with friends or family.

Second, even if one were convinced that agape was a normative ideal, it seems impossible. This charge of impossibility could take a number of forms. First, you might have metaphysical or conceptual doubts: what would love for all even be? We know what it is to love a friend or a parent or a child, but it is hard do get a grip on what it would be to love everyone—even those you haven't met. Or inasmuch as one can get a grip on it, it looks psychologically untenable for human beings. It's no accident that Buddhists think it will take many lives to achieve that ideal; many Christians likewise concede that the ideal of agape cannot be achieved in this life, but that it can be achieved only in the next life or in eschaton.<sup>7</sup>

And lastly, how could love serve as a foundation of ethics, not just an important piece but something closer to Paul's claim that the whole of the moral law can be captured by an injunction to love others? Granted that an ideal of love seems appropriate in some contexts; how could it play a foundational role in ethics as a whole, explaining the reasons we face and obligations we have in general.

Answering these three questions—why see agape as an ideal? how could agape be possible? and how can the ideal of agape serve as a foundation of ethics?—is the task of this paper. We'll begin with more ordinary love. I'll argue that when we attend carefully to the love we bear family, friends, and the like, we'll see that love has a kind of outward pressure: the reason for loving anyone is a reason to love all; or put another way, in our love for one, we can find an argument

---

<sup>7</sup>"The ethical demands made by Jesus [to love] are incapable of fulfillment in the present existence of man. They proceed from a transcendent and divine unity of essential reality, and their final fulfillment is possible only when God transmutes the present chaos of this world into its final unity" (Niebuhr, 1935, 56-57).

for the conclusion that correct response to the humanity of every human being is to love her. I'll then address worries about agape's conceptual and pragmatic impossibility, arguing first that while the most familiar cases of love are those we bear for particular people, entities or things, we also can (and do) have what I call "plural love," a love for some plurality of Fs. I'll argue that this kind of love is familiar (quite literally in that it is the kind of love many people bear for the members of their family) and provides the resources to explain what it would be to love all. But the threat of agape's impossibility will still loom large. I'll argue that what is impossible in practice is loving all of humanity while still acting on our love for individuals like family and friends. Most of us should not manifest agape. But I argue that we should not give up on agape's relevance in determining what to do; we should see it as an ideal that we must approach as much as possible. Though for those who do love family and friends, full agape is impossible, we can nevertheless approximate it. And in this idea, viz. that we should approximate the ideal of agape, we can find a novel account of the foundations of ethics. Where most agapist philosophers have defended a kind of consequentialism, I'll argue that what it is to approximate agape just is to manifest an attitude of respect for all. The hallmarks of an ethic of respect—universality, impartiality, an anti-aggregationist commitment to the separateness of persons, recognition of constraints against promoting individuals' welfare, and some kind of distinction between perfect and imperfect duties—can all be derived from agape. Respect is an approximation, a shadow, of love for all. Though we may not have known it, much of ethics—at least for those of a broadly deontological persuasion—has been done in the love's shadow. Appreciation of this fact paves the

way for a more foundation of ethics that is both more explanatorily unified and better speaks to the felt impossibility of doing right by the humanity of others.

## 1.2 A quasi-account of love: feeling and doing

To begin, we need an at least rough grasp of what love is. My hope is that most of my arguments do not rest on any particular theory of love. We do, however, need to give an at least rough characterization of some of love's features.

Of course, one of love's most striking features is its sheer variety. I love my parents, my partner, my friends, dogs, cats, Netflix, Mahler's 9th, and my favorite frying pan. Though we use the same word to describe our attitude in each of these cases, it is clear that the attitudes involved are importantly different. Even taking a kind of love that on its face admits of more unity—say romance or friendship—the diversity of love threatens to scuttle any attempt to say something about love in general.

But there is, I submit, much that unites all these attitudes and about which we can say a great deal. To make our task somewhat simpler, I will set aside love for inanimate objects and the like and focus on love for animals—and humans in particular.<sup>8</sup> Without pretending to offer a full analysis or definition, we can identify two constitutive features of love. The first is a way of feeling, viz. an emotional susceptibility to the beloved. The second is a way of acting; one who

---

<sup>8</sup>Briefly, “love” for an inanimate object is sometimes attributed when we like some object very much—“I don't just like ice-cream, I love it.” That we use the word love here signifies little more than intensity of feeling. But some love for inanimate objects is, I think, more unified with love for animals and people when a thing is valued in and of itself, though this is not the place to explore these issues further.

loves notices the beloved and takes her interests as especially weighty reasons for action.

To begin with the obvious, love involves a certain kind of feeling. It is commonly assumed that this feeling is a desire for the beloved—to be with him or to have some kind of connection with him.<sup>9</sup> But to take this to be essential to love is to take far too narrow a view—while many loves do involve such a desire, others do not. For better and for worse, many of us have friends or family who drive us crazy, whose company we would avoid but that we suffer out of love. Following a divorce, a divorcée can still love her ex-spouse while never wanting to see him again; she can care about him, wish him well, hold a place for him in her heart all while never wanting him in her presence. While neither love features a desire for company, it is nevertheless involves some kind of feeling—some tenderness of the heart.

Love involves a vulnerability to the other (Velleman, 1999). To love another is also to feel her happiness, her goods as well as her pain and the ills of her life; and it is to do so in a distinctive manner. Imagine reading a list of casualties of a terrible accident and seeing a long list of strangers' names. However saddening that might be, there would be a world of difference if on that list you found a name of one you loved: a family member, a friend. Likewise, the successes and welfare of our loved ones move us more like our own successes than do the successes and welfare of strangers; imagine the list again but as award recipients. The manner of feeling is akin to how we feel about the goods and ills in our own lives. We in some sense identify with those we love; we bring them into our hearts, their welfare, suffering, successes, needs, wants, etc...become ours.<sup>10</sup> As Nozick puts it,

---

<sup>9</sup>As in Jeske (2008) or Sidgwick (1981, p. 244) .

<sup>10</sup>See Korgaard's (1996) and (2009). Johnston (2010) takes that identification

“When something bad happens to one you love,...something bad also happens to you.... If a loved one is hurt or disgraced, you are hurt; if something wonderful happens to her, you feel better off” (Nozick, 1989, p. 58) We are emotionally susceptible to those we love. These metaphors—opening or connecting the heart, bringing another into the heart—capture the sense of emotional vulnerability. And it is in this vulnerability that lies much of the beauty and challenge of loving others.<sup>11</sup>

While these feelings are central to love, they do not exhaust it. Imagine someone who felt in this way distinctive of love but was never moved to do anything about it. Suppose I claim to delight in the happiness and feel the pain of another but when presented with an opportunity to make her life better—to help her achieve her goals, to make her happy, to save her from some harm—I am completely unmoved. It’s not that I am more strongly moved by some countervailing consideration, e.g., that interference would be objectionable in some way. I just don’t take anything about her life as a reason for action. I’m tempted to think that the case I’ve described is impossible or at least possible only for an agent with a severely fragmented psyche—our feelings and motivations are intimately connected, though the nature of that connection is complicated and contested.<sup>12</sup> In any case, I think it is clear that the motivationally impotent state is at best a defective kind of love; if someone claimed to love you and was so un-

---

quite literally.

<sup>11</sup>Of course, this can happen to a greater or lesser degree; the depths to which we open our hearts to our colleagues, to our friends, to our cousins, to our romantic partners, etc..., varies. That variance is quite significant, and as I argue in both subsequent chapters has implications for many of our obligations and permissions, including, crucially, our obligations and permissions surrounding truth-telling and deception.

<sup>12</sup>Cf. Ebels-Duggan (2009, p. 143).



moved to act, it would seem insincere or confused. Love involves a kind of motivation—to be moved by the welfare and preferences of the beloved qua reasons. What Kant called its “practical” as opposed to merely “pathological” component (1998a, AK 4:399).

While I won’t argue for this here, we can best understand the practical component of love as a habit of character, a principle of practical reasoning, or a way of looking onto the world that involves noticing certain facts and according them weight in one’s practical deliberation.<sup>13</sup> The courageous agent notices when others are in danger and her motivations are sensitive to that; compared to the rest of us, she weighs danger to others more strongly and to herself less so. So it is with love. Insofar as we love another, we notice him; his interests strike us as practically significant, and we give them comparatively greater weight. (What interests? Does one who loves respond solely to facts bearing on the beloved’s wellbeing or do things like the beloved’s choices, ends and projects matter as well? We’ll consider this below (§1.5.2). For now, we need only the idea that loving someone involves reasoning in such a way that takes the reasons stemming from them as especially weighty and important.)

---

<sup>13</sup>For others who take love to involve a some principle of practical reasoning or habit of character, see Pettit (1997); Kolodny (2003); Frankfurt (2004); Ebels-Duggan (2009); Spaemann (2012); Keller (2013); Setiya (2014). We can be quite neutral on what in our moral psychology a habit of character of practical principle involves. One might think of them as dispositions to be moved by certain facts (Setiya, 2010), or essentially rational dispositions to be so moved (Wedgwood, 2006; Arpaly and Schroeder, 2015). They could involve the adoption of a maxim and issuance self-legislation (Kant, 1998a; Korsgaard, 2009), a decision to treat something as a reason (Bratman, 1999a), or the setting of a relevant end (Bratman, 1999a). Thus, it might be that to love a friend involves a disposition to respond to facts about them, or something more: seeing yourself as a friend, seeing your friends’ needs under a distinctive normative guise (e.g., as reasons for you), deciding to take your friends needs as reasons, or setting your friends’ wellbeing as an end.

### 1.3 Love's Outward Pressure

I claim that we should love everyone. But love seems like a way of singling out a particular person vis-à-vis others as special. It involves attending to their particularities and valuing them in an exceptional way (Jollimore, 2011). We care about and do far more for our loved ones than we do for strangers. I notice when my friend Kevin is feeling a little sick whereas I often don't notice the fact that many strangers are dying. I don't attend to their far more dire needs. I value my friends so much that even knowing a \$50 birthday gift could pay for the deworming of 150 children or malaria nets for 10 families—serving strangers' interests far more than the gift serves Kevin's, I still spend the money on him. When even minor goods and ills befall Kevin, I feel them; when disasters strike around the world, I often don't. He is, to put it somewhat tritely, special. To love someone is to be partial towards them.

Moreover, part of what is so attractive about being loved can be this exact feeling of being singled out: my friends care about me in a way they don't others; my romantic partner attends to me in a way she doesn't others. And the language of love often reflects this kind of comparative exclusivity. We tell those we love that they matter to us in a way that others don't, that they are "special," that they are "the one," that they "mean the world" to us in a way that others don't. Even love's pathologies are a reflection of love's prioritization of the beloved over others. Where healthy love for friends, families and partners exhibits some partiality, pathological love can involve focusing on the beloved to the near total exclusion of the rest of the world. In short, love seems like the kind of attitude that is at home in the special relationships we bear to the few. Why think love should

be extended to all?

If we look closely at our love for individuals, however, we find a surprising result: there is a kind of outward pressure to expand one's love. In our love for one we find reason to love all. I mean this in two ways. The first turns on what it is that is good about loving others, viz. that what is good is better the more people we love. The second is based on why we love; the reason for the best kind of love is a reason that applies equally to everyone.

First, I take as a starting premise that love is an essential part of a good life. One who fails to have and to manifest love is missing something. But why is this so?

For one thing, it facilitates relationships that are important. You can't have friendship without love—or at least you can't have really worthwhile friendships without love—and so love is a necessary means to something important. But love isn't an essential part of a good life merely because it is a means to something else, viz. loving relationships. Something about the love itself is significant—for that matter something about the love is what explains why friendship is so valuable.

Love is a part of a good life in part because it is an expansion of the heart. With apologies for the flowery language, love is a transcendence of the bounds of the self. In being vulnerable to others and taking their interests as reasons—in bringing them into our hearts—we orient ourselves around something bigger than ourselves. We identify with others. Love involves, after all, a character trait. And is the sum of these character traits that make us who we are as agents. To love another is not just to feel something or to do something, it is to be something (i.e., a lover of that person), and that being is part of what

makes us the people that we are. To love, then, is to have one's being encompass another; it is to some degree to incorporate another being of inestimable value into one's self (Wolf, 2010). To live a good life means involves orienting oneself around and engaging with genuine value. Few things (if any) have more potential to bestow meaning than the inestimable value of another human being; and no form of engagement with a person is as deep or profound as love. To love another, then, is to engage as fully as possible with that which is the highest value. And that seems to be what is good about love—it is an orientation of a life around something genuinely valuable in a way that gives one's life meaning.

But if expanding one's heart by loving one is good, expanding one's heart by loving two is better still. Of course, I do not mean that if having a loving relationship is good then having two of those relationships is better. Two romantic relationships are not necessarily better than one, especially if part of what the partners involve want out of their relationship is the exclusivity of monogamy. But that the love which forms the backbone of a loving relationship is good does suggest that some other expression of love, like friendship, would still further enhance a life. And that accords with our ordinary intuitions about lives well-lived. We celebrate lives that are filled with love—with friendship, family, etc.... Someone with just one or just two meaningful connections in her life seems to be missing out on something. To love more people is to imbue a life with still more meaning, and so to the extent that we can, we should. That which is good about love admits of expansion; so we have some reason to be more expansive lovers.

But an obvious objection is on the horizon. Just because loving

one person is better than none or loving three or four is better than one, it doesn't follow that loving twenty or fifty people is better than seven or eight! Love is a difficult business. With expansion of the heart come real challenges and costs, with respect to both feeling and doing. If loving means taking the beloved's interests as reasons, loving more people will likely lead to less time, effort and attention being spent elsewhere. And being emotionally open to others is draining, even in the best of circumstances, to say nothing of the pain and suffering that love for another can bring. To love is to be open to others' pain, and that is an important reason to limit our loves. These can all be countervailing reasons not to love.

Granting that what is good about love admits of expansion—inasmuch as some love is good, more love is better. There may still be limits as to how much one should love. Indeed, a central task of this paper is to understand those limits and what should be done in light of them. But just because we have (very good!) reasons not to expand our love too far, it does not follow that we do not have some reason to expand our hearts still further. That is, acknowledging all of these important reasons to limit our love does not tell against the conclusion I'm defending: that what it is that is good about love gives us reason to expand our love still further. That reason may well be overridden, but it is there.

Imagine we were more radically expansive beings. (“Wait!” one might object! “Isn't this conceding that agape will be impractical?” Hold that thought.) Suppose that we had the ability to act on love for a great number of people without thereby running out of time, space and resources; imagine that we had the capacity to weather the emotional turbulence of loving many many people or that the world

contained little suffering with which to lovingly empathize. There, it is clearer that we should expand our love. This is to say that the reasons we should not expand our love have to do with those costs and challenges. Had we but world enough and time, we should love more; we don't, but that suggests that we should understand this as an outward pressure to love more that is overridden by reasons to limit that expansion instead of thinking that we only have reason to love a few.

This argument, however, shows us only that there is something good about loving, i.e., that we have reason to be one who loves expansively. It does not supply us with the reason for which we should love a person; after all, to love another in order to expand one's own heart would be oddly (and perhaps objectionably) self-centered. So while this argument is useful in giving us, qua theorists, reason to think that love of all is ideal in some respect, it does not yet tell us, qua agents, what reason we have to actually love anyone—let alone everyone. Famously, the reasons we have to be in a state or attitude of some kind may not straightforwardly transmit as reasons for forming that attitude as they are the “wrong kind of reason.”<sup>14</sup> But we can turn to a second argument that turns on the reasons for which we love, revealing a pressure that, as it were, speaks directly to the loving agent.

There is a significant debate in the literature over what, if anything, are our reasons for love.<sup>15</sup> I don't have anything further to add to it here, so I will be brief. I'll start by taking for granted that unconditional love is rational. To love someone unconditionally is to

---

<sup>14</sup>See, for instance, D'Arms and Jacobson (2000a,b); Way (2012).

<sup>15</sup>See Pettit (1997); Velleman (1999); Kolodny (2003); Scheffler (1997, 2010); Frankfurt (2004); Jollimore (2011, 2017); Setiya (2014) and my second chapter.

love him come what may, regardless of what changes he may undergo or what he may do. Such love is certainly not required of us, and we may have good reasons to love some people only conditionally (like those listed above). But it is not a mistake to love someone unconditionally.<sup>16</sup> A father might love his daughter even though she has come to treat him terribly, declares she hates him, and leads a life in which she harms others. He might love her even if he were to find out that she wasn't, after all, his daughter. His love for her could be such that whatever he learns, however she changes, and whatever she does, his love endures. What could be the basis of such a love that is not dependent on the beloved's being any particular way or having any particular relationship to the lover? Nothing but love on the basis of her humanity can make sense of that (Velleman, 1999; Setiya, 2014).<sup>17</sup> No other account (e.g., a qualities view on which we love others on the basis of their merits) can make sense of this kind of love.<sup>18</sup>

---

<sup>16</sup>More precisely, it is never a mistake to love someone unconditionally in virtue of the love's being unconditional. As I'll argue later, it is sometimes a mistake to love another, viz. when so doing would compromise the love you already have (see §1.6.2). But that has nothing to do with the love's being unconditional as opposed to conditional. Compare: "it is never a mistake to drive on the proper side of the road." Of course, if one is intoxicated, it would be a mistake to drive, but not in virtue of being on the proper side of the road. So in that sense, it is never a mistake to have one's driving to have the quality of being on the proper side of the road.

<sup>17</sup>Cf. Kierkegaard (1995, p. 86-89), which may endorse a view like this.

<sup>18</sup>Jollimore (2011) and his (MS) offer careful arguments that a qualities view can accommodate such love by basing the love in the beloved's historical properties. But such a view fails to account for the possibility that love could endure even a bout of severe amnesia; if the father wakes up and can't remember anything about his daughter but nevertheless loves her, that would not be a mistake Setiya (2014). Moreover, the historical account of the basis of love gets the phenomenology wrong. The father might love the daughter without it seeming to him like it is grounded how she was in the past. He could love her as she is, in spite of her flaws, and not just because of who she was, without thereby loving irrationally. That all being said, on a qualities view, part of what it is to love someone is to be especially attuned to whatever attractive features a person has; so while it is not as forceful, even on this view, we can find a kind of outward pressure. If I love my partner because she is beautiful, funny, kind, etc..., then, unlike Jollimore, I

Unconditional love is not merely rational, however. It is an ideal kind of love. It is no accident that traditional, Western wedding vows profess an explicit commitment to unconditional love—“for better or for worse, in sickness and in health, ‘til death do us part.” Unconditionality seems like something worth striving for. And it is not just romantic commitments that seem to be evaluable against a standard of unconditional love. For any love, of a friend, of family, it seems that the closer it is to being unconditional, the more deep and pure the love.<sup>19</sup> Consider two friendships; where one survives the friends moving apart from one another, the other doesn’t. The love of the former friendship seems superior qua love, and that is because it was less conditional.<sup>20</sup> Indeed, when we love, we often should deepen our love—we seem to have reason to approach and even achieve unconditionality. I don’t mean to romanticize it—it is not always the love that is called for, and when things do go awry, it can be trying. But as above, the presence of reasons against is no reason to reject that unconditional love is nevertheless an ideal kind of love.

And here, again, we encounter the outward pressure. To love

---

think we should see that as generating an outward pressure to love all who are beautiful, funny or kind. That doesn’t mean we are obliged to so love—see below. Moreover, Jollimore’s view incorporates the idea that love involves a commitment to seek out and be open to discovering new, lovable qualities in those we love. That is, part of what it is to love is to recognize there are reasons for love of which one is not yet aware but which one should seek. On pain of thinking our beloved is objectively exceptional, it is reasonable to assume (virtually) everyone has attractive qualities that, if we knew of them. Jollimore thinks the kind of attention directed at the beloved is essentially exclusive, but even on his view, what it is to love one puts one in a situation to see that there are reasons to love all (or at least most).

<sup>19</sup>cf. Kierkegaard (1995, p. 66ff.), which argues that only wholly unconditional love is purely love because, in contrast to conditional love, it is defined only by its nature as love and not by the particularities of its object. (That can’t quite be right, as unconditional love is still based on something, namely the humanity of the beloved.) See also Outka (1972, p. 17).

<sup>20</sup>Cf. Kierkegaard (1995, p. 78), which takes any form of conditionality to be a flaw in a love.



another is to respond to her humanity in a way that is not conditional on her being any particular way. From the fact that such love is rational, it follows that another's humanity is sufficient to rationalize love. But if that's so, then so, too, is anyone's humanity. To love one person unconditionally is to see that Freud's claim that "not all men are worthy of love" cannot be right (2010, 82). From the fact that unconditional love is love's ideal form it follows that in loving one, we are pressured towards loving on the basis of humanity. And if humanity is a reason to love, we have reason to love all human beings. Using "agape" to mean unconditional love, Evans puts the point exactly: "a person is a being such that 'agape' is the appropriate attitude" (1969, p. 134, emphasis original).

It is here that many object. The thought that we have reason to love everyone is sometimes treated as absurd on its face. Others offer arguments: to think we have reason to love all would make promiscuity rationally required (Jollimore, 2017). But in the first place it need not (and will not on this account) follow from the fact that we have reason to love all that we must do so. And whatever is objectionable about promiscuity isn't a matter of merely loving too many people; what's objectionable is breaking one's loving commitments or neglecting the loves one already has—and I'll have much to say on that in §1.6.1. Others say that loving all is just impossible given our limitations as human beings (Velleman, 1999), an objection the next two sections will answer. Those who have nevertheless accepted that we have reasons to love appeal to the notion of a special kind of reason (sometimes called "noninsistent") that justifies love but does not count as a *pro tanto* reason in favor of it.<sup>21</sup> On such a view, the

---

<sup>21</sup>See Jollimore (2011), who takes of reasons that make objects "eligible" but not "mandatory" for choice. Kolodny (2003) and Setiya (2014) both speak of

humanity of another is a reason to love but unlike ordinary (“insistent”) reasons, which can be overridden but always count in favor of doing something, it can be totally neglected without irrationality.

But in the first place, the motivations for appealing to this notion of a noninsistent reason are unwarranted. As the rest of the paper will show, we can accept the conclusion that the humanity of each is an ordinary (insistent) reason to love him while agreeing with the obvious fact that we are not required to love everyone. And as we’ll see (in §1.7), once the full view is on the table, treating the humanity of all as (insistent) reasons to love—i.e., as a reason to love of which we must take account though not necessarily, given counterveiling reasons, satisfy fully—can explain much the alternative view cannot.

We have, then, two arguments for the thesis that love has a kind of outward pressure, i.e., that we have reason to love all. But what if we are to meet the objections to such a conclusion, we must make sense of what this outward pressure is towards, i.e., of what agape could be.

## 1.4 Plural Love

What would a maximally expansive notion of love look like? Is the notion of loving all coherent? Even if it is, is it nevertheless too distant a possibility to make sense as an ideal for ordinary human life? One worry is conceptual or metaphysical: what would love for all even be? If love involves seeing the beloved as special, isn’t it impossible to

---

insistent and noninsistent reasons, notions taken from Kagan (1991, p. 378-81)—who after characterizing such reasons, goes on to say “I do not know what reasons of this kind would be like. The sort of considerations which are familiar from typical practical deliberation do not seem to me to be of this noninsistent nature. More than this: it is hard to imagine how reasons for acting could be of this kind (ibid., p. 379-80).

love all? If everyone is special, after all, no one is. Another worry is practical: to love another, it seems one must know her, and we can't know everyone. So while coherent, a love for all, so understood, seems practically impossible and so irrelevant.

If agape is to play a meaningful role in ethics, we need a conception of it on which it is possible for a human to manifest such love. The solution lies in seeing that in addition to loving particular entities like people, we can also exhibit what I'll call "plural love," a love for the Fs, for some F.<sup>22</sup> Agape should be understood as plural love for human beings; so conceived, love for all is coherent and practically realizable.

#### 1.4.1 Particular love and plural love

Love is, of course, for something; love has an object. The most familiar, paradigm cases of love are those found in relationships like friendship, romance and family, where the object of love is a particular person. It is easy to see how our two main features of love—the practical and emotional component—would take shape in love for a particular person. Love involves being emotionally vulnerable to that person and taking her interests as relatively weighty reasons to act.

We can also love social groups—particular entities that are made up of many individuals. I might, for instance, love a team, my department or my country; I'd be emotionally vulnerable to it, and I'd take its interests as reasons to act. I might, for instance, respond to the fact that buying Red Sox tickets supports the team and buy a

---

<sup>22</sup>I am here drawing on the notion of plural quantification introduced by Boolos (1984). While I will occasionally refer to "a plurality" or "the plurality" of Fs, this is just for convenience and can always be paraphrased away in terms of the Fs (without mention of a or the plurality).

few tickets; to the fact that hitting a homerun would win the game by wishing it happens; and I might feel success and failures of the team, exhibiting a profound emotional vulnerability to the object of my love.

Call this kind of love, love for a particular entity—a person or a group—“particular love”. It would be easy to think that this is the only kind of love there is. And if it were, agape would be hard to make sense of—at least for human beings. To love someone or something in this way requires some kind of acquaintance with the particular entity. That acquaintance might be thin—via a photograph, a letter or a perhaps fleeting glance—but it must be there to some degree.<sup>23</sup> But agape would have to be a love without acquaintance—we cannot meet or keep track of all human beings. Perhaps God could love all human beings the way we love each of our friends, but we cannot. If agape is to make sense, we must understand love of all as something other than particular love for each human being.

I’ll now argue that another kind of love is possible: not love for some particular entity but rather a love for some plurality of Fs. Agape—the love that would be the maximal realization of love’s outward pressure—should be understood as love for human beings, not as love for humanity conceived as the group of human beings (or the species), nor as a love for each individual human.<sup>24</sup> It is the task of this section to get this kind of love in view. To do so, we’ll forget agape for a moment and consider a more familiar case (pun intended):

---

<sup>23</sup>See Setiya (MS).

<sup>24</sup>One technical note about the logic of plural love. In the following section, I’ll talk of love for the Spencers and then go on to talk about love for the Humans. Standardly, when qualifying over some things,  $xx$ , the variable “ $xx$ ” is taken to rigidly pick out those particular things that are among the  $xx$ . So when I say Jack loves the Spencers, I really mean that Jack has an attitude for  $xx$  such that any  $x$  that is F is among  $xx$ .

love for members of a family. I'll use that case to get love for the Fs, "plural love," in view; I'll then go on to argue that loving the Fs has a number of key features: it is less personal than a particular love for each F; it does not require knowing each F one loves; and it is impartial among the Fs. To foreshadow a bit, an approximation of love will share these features, and it is ultimately these features that will provide the basis of my claim that respect should be understood as an approximation of agape.

#### 1.4.2 Love for the Spencers

What follows is a phenomenology of a case of familial love. You may not love your family in the way that the protagonist of our vignette does. It may be that in reflecting on what follows you should imagine Jack loving some different plurality of Fs, e.g., the members of his local community, his department, or some such. All that is required for the argument to succeed is that the love as described be possible, though I think it is moreover a love that is actual and at least in varying degrees not uncommon. The form of the argument will be a description of a case followed by (failed) attempts to make sense of the case in terms of particular love. The only way we can understand the case is by countenancing the existence of another form of love.

First the case: suppose Jack Spencer is attending a family reunion. The timing is a bit inconvenient and he'll have to make sacrifices to attend; furthermore, he knows time spent at the reunion could instead put to use in a way that served more people's interests to a greater degree, e.g., by volunteering for a charity or political group or simply making more money and donating it; and he knows that the money spent on flights, lodging and the toys he bought for all the young

Spencer children—many of whom he has never met—could have saved a life (or maybe even multiple lives) that will otherwise be lost. Why, in light of all that, would he go to the family reunion? It's no great mystery: because he loves the Spencers; he inconveniences himself, spends significant sums of money, and allocates his time out of love for them.

In what does Jack's love for the Spencers consist? First we'll consider two attempts to make sense of Jack's love in terms of particular love. Perhaps his love for the Spencers is simply a particular love for a many-person entity, viz., the Spencer family. Another thought is that it is particular love for each Spencer (i.e., love for Jane, and love for Gary, and love for Carmen, etc...). I'll argue that neither proposal can succeed and thereby show we need to understand his love as being an attitude not for some individual thing(s) but for them, the plurality of Spencers.

That a love for the Spencers can't be understood as a love for some special entity like the family is obvious. Jack might not have any thoughts about the family as such. His love might only be for people, not for a family. But even supposing he loves the family, that love would be different than a love for the Spencers. He might for instance love some things about his family, e.g., its size; but it would be confused to say he loved the Spencers for their size. It also seems consistent with a love for a family not to care very much about at least some members of the family. You could imagine a patriarch of an ancient family caring very much about the family—its reputation, its holdings, its power—while showing an utter disregard for most of the (unimportant) members of that family.<sup>25</sup> But we can suppose

---

<sup>25</sup>If I can be forgiven the pop culture reference: A Song of Ice and Fire's Tywin Lannister perfectly embodies the distinction between love for the Lannister family

Jack has a different kind of love; one that would be incompatible with disregard for any Spencer. That is, while love for some Fs entails a loving attitude toward each F, love for a group of which some Fs are a part does not; so we cannot understand Jack's love for the Spencers as (particular) love for the family itself.

One might instead think Jack's love for the Spencers should be understood as a sum of particular loves for each Spencer. But in the first place, such a view would not explain why Jack's love would extend to new Spencers automatically. It seems natural that if Jack loves the Spencers and a new Spencer is born, adopted or marries into the family, his love would embrace that new Spencer. But a love for a number of individuals would not explain that automatic extension.

Relatedly, the view could not explain how Jack's love can extend to those he does not yet know. The vignette stipulated that Jack bought toys for young Spencer children he hasn't met. And it is natural to explain why he did so in spite of the cost to him, and the vast good the money spent could have done for the distant needy, by appealing to simple fact it was done out of love. Everything he is doing about the reunion is done out of love. But that means that Jack can act out of love for those he has never met. Earlier I asserted that loving someone in particular required some acquaintance with them—some way of knowing him whereby he in particular can feature as an object of love. But where some would take that as a reason to revise the natural description of the case, I take it as evidence that Jack can act out of love that somehow relates to the children he hasn't met without featuring them in particular as objects.

---

(which he has in spades) and love for the Lannisters, which he lacks. He sees his children as tools to be used in the advancement of the family, not as objects of love and care.

Maybe, then, Jack's love for the Spencers should be seen as a sum of particular loves for each Spencer plus a disposition to love each Spencer Jack meets. That would at least go some of the way towards solving the problem of automatic extension to new Spencers, although it gets the phenomenology wrong. Imagine Jack meets Justin, who had become a Spencer by marriage some years back but had never had the chance to meet Jack. It would be natural for Jack to feel that he had already loved Justin and for Justin himself to feel already loved—that feeling can be one of the best parts about meeting family for the first time. The view under consideration, by contrast, would have it be that the love is formed in the meeting, not already present. Of course, a new particular love for Justin could be formed upon or after meeting him; a love for the Fs could even explain a disposition to get to know individual Fs and come to love them in particular. The point is that some love (i.e., the love for the Spencers) already extends to Justin.

And decisively, this revised view does not in any way explain the first problem of Jack's love towards people he hasn't met. Jack can already love those he is not in a(n) (epistemic) position to love particularly. He can buy toys out of love for the young Spencers. He could set up a college fund for the young Spencers out of profound love for the Spencers even while being unable to love them in particular.

In addition to love for a particular individual, we must recognize love for the Fs as another kind of love. In addition to particular entities, love can have pluralities as its object. Thus we can say that Jack loves the Spencer children inasmuch as he loves the Spencers, something that already seemed like a natural description of the case. Having given an argument that we need to countenance plural love,



I'll now characterize it.

The first thing to note is that it is less restrictive than particular love; it allows loving people one cannot love in particular. I mean this in two ways. The first is epistemic and that which we considered above: Jack cannot love unmet Spencers particularly, but his love for the Spencers (qua the Spencers) can extend to even to those he has not met.

Second, plural love can transcend psychological limitations. Notwithstanding the fact that we have reason to love everyone, some people are quite hard to love. They annoy us, bother us, or hurt us or those we love. Imagine Jack has an obnoxious uncle Gary; Gary treats others well enough, but whenever the conversation turns to politics, he expresses deeply bigoted views. But while Jack finds Gary almost intolerable, Jack loves Gary's wife and children very much. We all have Garys in our lives—and perhaps many are not in the end worth loving<sup>26</sup>; but suppose Jack wants to love Gary even though he finds Gary repulsive. Gary's individual personality is not, given Jack's psychology, lovable; but even if he can't bring himself to love Gary qua individual, he could love him as one among the Spencers. That is, he could form a love for some Fs (here the Spencers) that includes Gary just as his love for the Spencers includes the children he hasn't met. Jack's love for the Spencers is the same with respect to any particular Spencer; it is not a love for Gary qua particular individual (in all his obnoxious particularities). So just as Jack's love for the Spencers encompasses his sister, the kids, etc..., it can encompass Gary. So far as the love is concerned, he is just a Spencer. And, again for better or worse, I suspect that inasmuch as any of us do love the

---

<sup>26</sup>As I've emphasized, that we have a reason to love all does not mean that we have decisive reason to love all; we can have countervailing reasons.

Garys in our lives, it is just this way—by forming a love that does not respond to the beloved in his particularity but encompasses some plurality of individuals of which he is one.

Jack's love for the Spencers is thus less restrictive because it is insensitive to the features of each Spencer that make him or her unique. He need not know about them or if he does, they are shielded, as it were, from the point of view afforded by a love for the Spencers. This is all to say that love for the Spencers is a less personal kind of love; it is more expansive and less limited (epistemically and psychologically) because it is in this way thinner. It is not an improvement on particular love nor an adequate substitute for it; it is something different. Imagine Jack bore no particular love for his sister Jane and loved her only as one among the Spencers. She would be warranted in feeling hurt. We could imagine her complaint: "Your love for me is just the same as it is for all those cousins you've never met or for Gary! What about me? How could your love not be individualized to me—to reflect what makes me who I am or what has shaped our relationship and history together?" In this way, it is not a love suited to intimacy, which involves responding to and loving that which makes the beloved unique. This is not, to be clear, a criticism of plural love; Jack, after all, can bear a particular love for Jane alongside a plural love for the Spencers. I am just trying to bring out the less personal nature of plural love.

Along with its less personal character, plural love is also by nature impartial. Recall that love involves a practical component—a habit of character or principle of reasoning that takes the interests of the beloved as especially weighty reasons in deciding what to do. Love for the Fs involves a principle that is impartial among the Fs. It

is, of course, partial in one sense. Jack's love for the Spencers leads him to be partial to the Spencers vis-a-vis others; but among the Spencers, acting out of love for the Spencers would lead Jack to treat the Spencers equally. After all, nothing in the attitude could serve to pick out any one Spencer's interests and elevate them over others. If Jack were partial to one Spencer—if he took her needs and preferences as more important than that of any Spencer, if he attended more to her—then it must be in virtue of some attitude that is apt to pick out that particular Spencer, e.g., a particular love for her that supplements the plural love for the Spencers.

To sum up, love for some Fs is not reducible to particular love in some way—it is love for plural objects. Because it is a love whose objects appear not as individuals but as some among some Fs, it is a love with key features: it is not individualized to each F, it is less personal than particular love for a given F would be, and it is thereby able to transcend the epistemic and psychological limitations of love. And while love for some Fs involves being partial to them over others, it does not discriminate among the Fs and so is impartial among them.

## 1.5 Agape

With plural love in view, we can see what agape could be. It is the task of this section to describe it. I'll first argue that we should understand agape as plural love for human beings; thus conceived, it is possible for a human being to manifest. I'll then go on to draw out the implications of what an agapic agent, so understood, would be like. In the simplest terms, she is an agent who is emotionally susceptible to all human beings and who takes the interests of each

human being as a reason to act. I'll go on to argue that this concern for human beings is marked by five structural features, each of which can be derived from our quasi-account of love (§1.2) and the previous section's understanding of plural love. The agapic agent's love is:

1. Universal: The agapic agent's love is for all.
2. Impartial: The agapic agent treats all equally.
3. Anti-welfarist: The agapic agent is concerned with the welfare of others, but is likewise respectful of their choices. She reasons in accordance with constraints grounded in individuals' choices and will not promote another's wellbeing when doing so infringes on their autonomy.
4. Anti-aggregative: The agapic agent does not see individuals as fungible. Her reasoning is committed to the separateness of persons. In other words, the agapic agent reasons in accordance with certain constraints when making interpersonal comparisons.

In other words, the agapic agent has an impartial regard for all that is at once about their wellbeing and is concerned with honoring their autonomy as agents. That an agent is genuinely agapic entails that she respects all. Where this section will argue for that entailment, the next will go on to make a claim about explanatory priority: that the imperative to respect can be derived from an ideal of agape. But for now, our task will be to articulate the features of agape.

One note: because ultimate the aim of this paper is to argue that deontology's proper foundation can be found in love, I am trying to remain neutral with respect to various debates between deontological

views, e.g., exactly what anti-aggregative or anti-welfarist constraints there are. Are these to be understood as rights? Are they absolute or can they be overridden? Are there constraints against harm, lying, bodily trespass, use of others' property? And what is the exact shape of those constraints? These are all questions that can be raised in terms of the agapic agent, i.e., under what, if any circumstances, would an agapic agent harm another or use another's property without permission? And I think that much progress can be made in ethics by thinking through debates within deontology in terms of love.<sup>27</sup> But given that my purpose here is ultimately to make the foundational claim that respect should be explained in terms of love for all, my sketch of agape will be just that, a sketch of the general outlines of an agapic approach to the world but not its details.

### 1.5.1 Universal and Impartial Regard

We could not bear particular love for every human being, at least not in a world that has a moderate number of people in it. Love for all, then, must be understood as plural love for human beings. Such a love would encompass every human being, and automatically extending to every new human being that comes into existence.

(What about non-human animals? A Buddhist would object that agape, or whatever expansive love is really worth thinking about and regarding as an ethical ideal, is not a love for human beings but for all animals or all conscious beings or maybe even for all living things. But the arguments I offered earlier about the outward pressure of love don't obviously extend to animals or living things as such; if the basis

---

<sup>27</sup>Progress I try to demonstrate in my third chapter, which argues that our obligations around telling the truth can be understood (surprisingly) only in terms of loving relationships.

of my love for individual people is their humanity, not their animality or consciousness or life, the love that fully and non-accidentally responds to that basis is a love for human beings (or if you'd like, for people), not animals, the conscious or the living as such. Those categories are all even more expansive, and while the argument in favor of expanding one's heart may suggest we should love all animals or all life, it does not follow from our endorsement of unconditional love for human beings that we have a reason to love non-human animals or the living in general. There may be other reasons to think that we should love all animals; indeed, I think we should, and that ethical issues surrounding our treatment of animals can be better understood by thinking through a framework of love. But for now, I will set these issues to the side and restrict our attention to love of human beings.)

From what we have said about love in general and the nature of plural love, what can be said about an agent who manifests agape? In the first place, it is not just accidentally but essentially universal. Just as Jack's love for the Spencers essentially encompasses every new Spencer, including new ones, so, too, would agape essentially encompass every human being.<sup>28</sup> It is likewise essentially impartial. Because it is a love for human beings as such, there is nothing in the attitude to single out any particular person vis-a-vis any other. The attitude borne towards one is the very same attitude born towards every other, so the way in which it responds to each is essentially equal.<sup>29</sup>

---

<sup>28</sup>See Kierkegaard (1995, p. 58). Many other Christian authors agree (Outka, 1972, 9-24). A notable exception is Barth, who thinks God's command to love thy neighbor extends only to those within the community of Christians (Dogmatics, IV/2:802-4, as cited by Outka 1972, 210.)

<sup>29</sup>Again, this is a widely endorsed position by those working in the Christian tradition (ibid., 9-24; Niebuhr 1957, 108; Wolterstorff 2011, 21-40) and also by Gootman (2008, Lecture 1), who, like Niebuhr, makes much of the command to

What would impartial, universal love for all look like? As with any kind of love, it should be seen as a trait of character involving an affective and practical component. Let's take each in turn. The agapic agent is emotionally susceptible to all (and equally so). At first pass, that might seem psychologically impossible. Wouldn't our limited emotional bandwidth make universal vulnerability impossible (Velleman, 1999)? No. Consider ordinary, day to day love. Our love for all our friends involves emotional susceptibility to each of them; but it is consistent with our loving a friend that on a given day, we are emotionally exhausted, perhaps by love of others friends, and so cannot feel with him. Having taken on the stress, pain, joy, etc. of some of our beloveds, we can (temporarily) run out of emotional room for that of others. But we don't thereby stop loving them. Susceptibility is a disposition or propensity to feel with and for others and like any disposition can be masked. The point is not the the agapic agent would actually feel for all human beings; that would be impossible, at least for human beings. It is that she is (equally) vulnerable to all. As she moves through life, the sorrows and triumphs of others move her directly and deeply, as we are moved by only a select few. Where most of us have a(n important!) tendency to ignore most people and, when ignorance is impossible, harden our hearts to their troubles, the agapic agent does neither, letting it all in. An example to illustrate: suppose you are on your way to a close friend's party and you see an accident on the highway. Whatever you happen to glimpse, you would likely have an only limited emotional response to the pain of the victims. If the response were strong, you might even try to push those feelings out and harden your heart so as not to show

---

love the neighbor "as thyself."

up to your friend's party an emotional wreck who can't share in her joy because of some strangers' pain. Now suppose instead that it is Jack who sees the accident and he knows the car that crashed was full of Spencers; even if he were headed to a friend's party (and of course he'd likely stop! but put that aside), he'd have much more trouble pushing those feelings out—he's vulnerable to them. The agapic agent is like Jack but with respect to everyone. Such a life would be overwhelming and all consuming (more on this in section 1.6.1) but not impossible. And we may know people who, though they may not be fully agapic, nevertheless manifest some fraction of this tendency, i.e., those who have thin emotional walls and are very affected by those around them, even those they don't know personally.

Much the same can be said of the practical component. While the agapic agent may not be able to act out of love on behalf of every person, she can adopt a principle of reasoning that takes each person's interests as reasons to act. In some sense anyone who takes their moral obligations seriously sees all as equal (more on this below, §1.6.3), but in another we do not. Most of the time, we don't really take account of the needs and interests of anyone but a special few; we might allocate some (possibly great) portion of our time, money, effort and attention on the many, but when it comes to spending an evening with a friend who needs a little comforting, we do it despite (often without any thought for) the fact that an evening spend helping those in desperate need could do far more for them than our friend. While we always accord some minimal regard for others, we often don't take others into account any more than that. The agapic agent is in that respect quite different. She does not privilege any one human being; she is equally disposed to notice the needs, interests of each person;



she takes the interests of each person as (equal) reasons for action. Where we don't notice, she does; where we act on behalf of the few, she acts impartially.

## 1.5.2 Wellbeing and Autonomy

Does this mean that the agapic agent is consequentialist—concerned solely with promoting the good of all?<sup>30</sup> All of those in the seventeenth and eighteenth centuries who took a love-centric approach thought so,<sup>31</sup> and some contemporary authors on love have likewise thought its practical component consisted largely of taking their welfare as reasons (e.g., Frankfurt 2004). But agape is, in two crucial respects, non-consequentialist. First, the agapic agent cares about more than what is good for the beloved. She cares about his choices, ends and projects. Second, she would not think aggregatively about those she loves. Taking each in turn, let's first consider what kinds of things would motivate an agapic agent when considering just one person she

---

<sup>30</sup>By consequentialism, I mean the old fashioned kind, before attempts were made to “consequentialize” deontological constraints, yielding “consequentialist” views that align with deontological predictions, as in Brown (2011); Dreier (2011). For an overview and objections, see Muñoz (MS).

<sup>31</sup>For an overview, see Schneewind (1998). The first was Cumberland, who is often credited with first articulating something like a greatest happiness principle (1727, §1.iv; for a list of similar passages, Schneewind 1998, p. 103 n.6) which requires us to maximize “joint” (i.e., aggregated) happiness (1727, 1.xxxiii; see also 5.xix, 5.xxxv). Cambridge Platonists like Smith and More likewise took the injunction to love as a command to promote the “happiness of all mankind” (Smith 1660, p. 154-56; More 1666, p. 15, from which the quotation is taken). Hutcheson went so far as to define even rights in terms of the “general good” (1738, Vol. II §VII.VI, p. 297-98). Butler, who thought Hutchesonian benevolence was an important principle in ethics argued it could not be the sole exactly because it would have unacceptably consequentialist upshots: licensing the theft of another's property and making one indifferent to whether a good was had by a stranger or a loved one (1740, ¶8-10, p.383-4). Leibniz's love-based ethics were also consequentialist (1969, §4). A quick note that Hume saw an intimate connection between love and virtue, but not in that love was a virtue; rather, for a trait to be virtuous is “equivalent” with “the power of producing love or pride” (2007 3.3.1.3/SB575); see Korsgaard 1999.)

loves.

With respect to just one person, agape is similar to ordinary love for family and friends. So we can turn to particular love, and as I'll argue, we can see that that love is not an attitude that cares only about the welfare of the beloved (Ebels-Duggan, 2009). To love another individual well entails respecting that individual's autonomy. It is to take not only their wellbeing but also their choices, ends and projects as providing reasons for action. And it is to be committed to supporting them even when doing so would make the beloved worse off. In other words, love is by nature anti-paternalistic.<sup>32</sup>

These may strike some as odd claims to make. If ever paternalism (i.e., an overriding of another's autonomy for their own benefit) is appropriate, it is in the context of close loving relationships. Parents often should act in their child's interests whatever the child's preferences. We sometimes intervene in our friend's lives—we get in their business, trying to direct and effect their choices—in a way we shouldn't with others. Love and respect for a person have struck many as pulling us in opposite directions. Kant describes the “laws for human beings' external relations with one another” as a balance of “attraction and repulsion”:

The principle of mutual love admonishes them constantly to come closer to one another; that of respect they owe another, to keep themselves at a distance from another; and should one of these great moral forces fail, “then nothingness (immorality), with gaping throat would drink up the whole kingdom of (moral) beings like a drop of water.”

---

<sup>32</sup>There may be limits here—it may be that love would sometimes have us override such constraints against paternalism when enough is at stake. As above, I won't settle such issues here.

(2017, AK 6:449, *emph. original*, quoting the poet Haller)

Where respect for autonomy is in part a matter of respecting the difference between what's mine and yours, Kierkegaard claims that "there are no mine and yours in love" (Kierkegaard, 1995, p. 265, *emph original*), and he is clearly on to something right.<sup>33</sup> In short, it seems reasoning in the manner of one who loves and one who respects an individual are quite different. But I'll here argue that isn't so. Throughout, I'll make claims about what is more or less loving or what meets the standards of love. I don't mean to suggest that being insufficiently loving or failing to meet such standards means one's attitude isn't love; I am making a claim about what it is to love well, *i.e.*, what it is to respond fully to the reasons of love, in the same way we might talk about what it is to be more or less courageous or generous.<sup>34</sup>

In the first place, it seems that one of the most loving things we can do for those we love is support their own choices and ability to author their lives as they see fit (Darwall, 2018).<sup>35</sup> To support another isn't just to make them better off, it is to enable them to pursue their own ends; supporting a beloved means being sensitive to what it is they want to do. For a parent to help pay of student loans and thereby

---

<sup>33</sup>Nygren and Niebuhr likewise see a conflict between love and justice, and we can see how their remarks would extend to respect. Nygren takes forgiveness to be the paradigm expression of love, and understands forgiveness as something like the willful neglect of what is owed (1953, 88-90). And see Wolterstorff (2011, 51-55) for a decisive response. For Niebuhr, consideration of rights or what is owed only makes sense in a context of conflict and competing interests whereas love exists in a space of harmony (1957, 28)—hence his thought that *agape* must wait till the end of days when this world of strife is eliminated (1935, 56-57).

<sup>34</sup>Cf. Ebels-Duggan (2009). I also don't here make any claims about the grounds of those standards—they could be constitutive norms derived from the nature of love (Korsgaard, 1997, 2009), explained via a naturalistic reduction Setiya (2010), or fundamental normative truths about what reasons a loving agent should respond to.

<sup>35</sup>This is what Darwall calls love's "upholding" of the beloved (2018, pp. 105-6).

enabling her child to pursue what he would is a paradigmatically loving thing to do; for her to force him into the career she thinks best seems unloving.

In general, actions which disrespect those we love seem not merely disrespectful, but unloving. That is, actions which fall short of a standard of respect for another's autonomy at the same time fall short of a standard of being fully loving. Suppose you know your friend is dating someone (or pursuing a career, or making some other life choice) that is bad for them; that is, you correctly believe that what they are doing will make them worse off, but even when you outline the consequences, your friend is firm in their choice. Suppose further that you correctly believe that with just the right threat ("It's him or me; if you don't stop seeing this guy, I'll stop speaking with you") or just the right lie ("I saw your new boyfriend with another man; he told me he'd just lie to you if I said something, but he's cheating") you could in fact make his life better off. Such actions would be wrong. (If your views require it, adjust the degree to which the lie or threat would make the beloved better off to secure this verdict.) But in addition, they strike us as unloving. We might, of course, try to justify such actions by appealing to our love for the lied to or coerced person ("I'm doing this because I love you"). But such appeals express a confused view about love. That disrespect towards those we love is also unloving suggests that love for an individual requires respecting him; actions towards one's beloved which fall short by the standard of respect also fall short of a standard of love.

And (by way of a second argument) that makes sense. To love someone requires seeing and responding to the other as what he is. And human beings are not merely receptacles of wellbeing; they are

beings with a capacity for choice. To treat someone in a way that disrespects their autonomy, even for the sake of their wellbeing, is to see and respond to only a part of him, and in that way is deficient as love.

I take these relatively simple considerations to give us defeasible reason to accept that loving another requires respecting him. The real work of defending that claim, however, is responding to what seem like powerful lines of argument to the contrary. Respect involves letting another person live his own life; but it is with those we love that we seem to have the most license to interfere. Perhaps we cannot lie to or coerce our friend out of a bad relationship or career, but we can certainly get up in their business and try to get them to see that they are making their lives worse off. If respect involves honoring the distinction between mine and yours—my property and your property, my business and your business, my life and your life—isn't Kierkegaard right to say that love is the elimination of that boundary? Of course he is right; but once we appreciate why, we can see that Kierkegaard's insight, and even the truism that we can and sometimes must override the wishes of those we love for their own good, is entirely consistent with the idea that love requires respect for another. In fact, we cannot even make sense of the idea of Kierkegaardian transcendence of the distinction between mine and yours without thinking in terms of autonomy and respect.

While interfering in the life of a friend or family member is one thing, interfering in the life of one who does not reciprocate one's love or with whom one has no relationship is another thing entirely. Suppose that after a divorce, one member of the once pair continued in his love for his ex while the other's love died; suppose moreover

that the divorce was difficult and the two have no ongoing, mutual relationship with one another, i.e., they aren't friends. If the one whose love survived continued to interfere in her ex-wife's life despite the latter's wanting nothing to do with her, she would be acting, like above, both wrongly and unlovingly. But the difference between such a case and one where an agent may or even should interfere isn't the love the agent bears. A spouse can interfere in her wife's life where a divorcée cannot because of the attendant circumstances, i.e., because of the relationships that they are in.

Loving another and being in a loving relationship with her are two different things. One can love unrequitedly; one cannot be married, or be friends with, or date unrequitedly. Indeed, it can even be that two people bear love for another without being in a relationship together, as in a case when two people love each other believing the love is unrequited. A relationship involves more, viz., a reciprocal recognition of the other and a mutual embrace of the each in the others' life Darwall (2018). And part of what is involved in coming to form loving relationships is the invitation of the other into one's life. We let our friends and partners in because we love them. Becoming friends makes it the case that our otherwise private life is now our friend's business, our property is free for them to use, our secrets our their secrets—at least up to a point. Much the same is true in familial relationships, although there the permissions are not always generated by invitation. Mothers, fathers and guardians are permitted to override their child's choices not because they love their child or because the child loves them back but because of their relationship. “That is why in adult relationships the accusation ‘you're treating me like a child’ is a potentially valid reproach” (Ebels-Duggan, 2009). Rela-

tionships have their own rational significance, generating permissions and obligations that merely loving another does not.

When we intervene in a beloved's life permissibly, it is not a violation of their autonomy; our relationship permits us to so intervene. And if our interference steps over the line into the areas of their life that our not ours to participate in, it then becomes disrespectful and unloving. Kierkegaard is right that love eliminates (some of) the distinction between mine and yours in that love for another engenders the formation of intimacy, closeness, and relationships in which we waive constraints that apply to others. That is, loving another does not eliminate the distinction between mine and yours or pull us in the opposite direction as respect. Love for another typically leads us to form relationships that change what is involved in respecting the other. Put another way, what it is to respect another's autonomy depends in large part on the kinds of relationships one is in with the other, with some relationships allowing more intimacy and interference—if it can even be called that—without thereby affronting the other's autonomy.<sup>36</sup> But that shows no conflict between love and respect. Instead of ignoring what respect requires, love often seeks to reshape it, making what was mine or yours instead that which is ours.

Thus, love involves taking as reasons not just the wellbeing of the beloved but also facts about their autonomy. To be a fully loving agent requires taking the choices of the beloved as likewise generating reasons to act and recognizing constraints against promoting her autonomy. That is all to say that the agapic agent will therefore be sensitive to individuals' autonomy.

---

<sup>36</sup>See my third chapter. Cf. Ripstein (2009); Pallikkathayil (2010).

### 1.5.3 Anti-Aggregation

Just as the agapic agent will see autonomy as constraining and guiding her choices with respect to an individual life, so, too, will it guide her reasoning about actions that touch many lives and involve the opportunity for trade offs between them.

To see this, first ordinary love of friends. Imagine Sophie has two friends, Jan and Eva, whom she purports to love. Sophie, however, constantly lies to and harms Jan for the sake of Eva. If Jan has a donut and Sophie (correctly) believes that Eva would be made better off by the gain of the donut than Jan would be made worse by its loss, she secretly steals it and delivers it to Eva. If a limited-ticket concert that Eva would enjoy more is close to selling out, she lies to Jan, saying the tickets are gone, to increase Eva's chance of getting one. Such an attitude towards Jan seems inconsistent with properly loving him. Indeed, if when caught, Sophie explained "I was just doing what was best overall for my friends—had Jan enjoyed donuts or concerts more, I would have done the opposite," both Jan and Eva could complain that Sophie didn't really love them or didn't love them well. To love individuals is not to take a consequentialist attitude of beneficence towards them—love recognizes interpersonal constraints.

Why? Because to love another is to see her as not fungible. It is to see the beloved and her wellbeing as not simply substitutable with that of another. One who loves well would not feel that benefits to one beloved compensate for harms to another. And just as we are with particular love for friends, the agapic agent would not see those she loves as fungible either. Agape's character as a plural love guarantees this. It is a love for them, not for some it; by its nature, it treats its objects as a plurality of individuals. One cannot love the



For without recognizing that they are many individuals. In this way, plural love contrasts to particular love for a group and is like bearing particular love born for many individuals. For instance, one who loved humanity the species, not the plurality of human beings, would see the individuals as fungible. That would be a love that related to individual humans as part of a larger whole, whose relevance to the lover was as pieces that contributed to the object of love. But plural love is by its nature committed to the separateness of its objects.

#### 1.5.4 Summing Up

We've seen that agape entails respect. To love one requires respecting his autonomy. Therefore, the agapic agent takes not just the welfare but the choices, projects and ends of others as providing reasons for actions. She sees constraints against promoting an individual's welfare against his will and reasons in a such a way that honors the separateness of persons. Moreover, her love is universal and impartial. But to love all far outstrips merely respecting all. The agapic agent takes the interests of all as equal reasons for action. Where it is not a failure of respect to, say, allocate some sizable portion of one's life to the needs of strangers but live the rest of one's life without giving much thought to the distant needy, the agapic agent does just that and is moreover emotionally susceptible to all, even those to whom the rest of us (permissibly) harden our hearts.

### 1.6 Agape as an ideal

So much for what agape is. What relevance does it have for us? Hopefully, I've described it in such a way that makes the idea of agape

as a practical ideal intuitive, i.e., an action or manner of reasoning could be held up against agape and sensibly judged by how well it approaches agape. But even if I have succeeded in showing a love for all human beings is in some sense possible, it still seems exceedingly remote. Maybe for some agents, in particular those without any special relationships, agape is a real option for their lives. But such an orientation to others seems unrealistically out of reach for most of us. To live a life of agape would require, among other things, giving up much of what matters in our lives, not least the particular love we bear for friends, family and the like. So what, if any, normative relevance could agape have for us?

In this section, I'll argue that agape is domineering, crowding out much else from the life of the agapic agent. We should therefore not regard agape as a requirement, nor even think of it as supererogatory—i.e., that which we should manifest but are not required to. Instead, I'll argue that we should approximate agape consistent with our other loves. A view that holds agape is to be approximated will, surprisingly, turn out to look just like familiar deontology.

### 1.6.1 Agape dominates a life

Loving another is a serious business. To act on the interests of some leaves less room to do so for others; to feel for some can make it harder to feel for others. We have finite resources—emotional capacities, time, attention, effort—and so, too, is our capacity to manifest our love.

Consider first what it is like to love someone in dire straits. Imagine Keerthi is a mother whose son, Adnan, is fighting a painful, and thus far losing, battle with leukemia. She shares his pain and helps

him however she can, day and night. Supposing the (significant) needs of her son exhaust her time, energy, etc..., Keerthi will be less able to act out of love for others. She just won't be as able to manifest her love for them. Emotionally things are much the same: it would be hard if not (psychologically) impossible for her to feel with others given the enormity of her son's struggle. Her emotional capacity would likewise be exhausted. Does it follow that Keerthi cannot, therefore, love others while loveing her dying son? If love were a matter of actually doing for or feeling with others, the answer threatens to be "yes." But there is a distinction between loving another and manifesting that love. Keerthi could still be emotionally susceptible to others and take their interests as reasons, but her ability to manifest her love would be impaired, both practically and emotionally. Given Adnan's suffering, Keerthi's love for him crowds out other loves from her life, not in extinguishing her other loves, but leaving no room for them to play an active role in guiding Keerthi's life.

Of course, ordinary love doesn't always or even typically consume a life like that. It depends on the circumstances of the beloved and the depth of the love. Agape, by contrast, is almost inevitably so consuming as there just are so many people in situations like Adnan whom the agapic agent loves.<sup>37</sup> Given the (contingent but firmly entrenched) realities of our world, the agapic agent is always in a position to save lives that will otherwise be lost, comfort those who will otherwise be alone, etc.... She will always be emotionally consumed by others since there are so many people she loves in such dire straits. To love all is, practically speaking, to guarantee that one's resources—

---

<sup>37</sup>cf. Nygren, whom Wolsterhoff describes as claiming (without argument and never explicitly) that agape is "jealous; it tolerates no other forms" of love (Wolsterhoff, 2011, p.30).

time, effort emotional capacities, effort, etc.—will always be at their limit. In the same way that Keerthi, consumed with thoughts of her dying son, would have little left to give to other friends, an agapic agent would have little to no space for those who are not in desperate need. After all, if one you love is dying or suffering greatly and you can do something about it, you'd be far more moved to help them than to spend time with friends who are doing fine. That others are likewise in a position to help is irrelevant. Imagine someone telling Keerthi that someone else could comfort her son and so she should come out for drinks with friends; that someone else could means nothing if her son goes uncomforted. So the fact that many people can save those whose lives are most at risk would mean nothing to the agapic agent if the lives are still being ignored.

The agapic agent is virtually guaranteed to neglect much of what makes life valuable in acting out of her love for all. As Cottingham puts it,

A world in which I accorded everyone at large the same sort of consideration which I accord to myself, my children and my friends would not be 'one big happy family'; it would be a world in which affection no longer existed because the sense of 'specialness' had been eliminated. It would be a world where much of what gives human life preciousness and significance had disappeared. (Cottingham, 1983, p. 90).

And notwithstanding my arguments that the agapic agent can, like Keerthi, still bear particular love for individuals, those individuals won't feel that love. "To an ordinary human being, love means nothing if it does not mean loving some people more than others" (Orwell,

1949, p. 1353), and the special relationships grounded on that love, that are sustained by individuals' partiality towards each other, would be impossible for an agapic agent.

It's no accident that those who devote themselves to the religious traditions that take love for all as an organizing ideal often forsake particular love—from Siddhartha Gautama leaving his family to St. Benedict's code forbidding friendships (Spaemann, 2012). There are other reasons besides agape's domineering nature, of course; but the tension between devoting oneself wholly to love of all and loving a child or partner is reflected in these traditions. Likewise is it reflected in the lives of those who seem to live out something like agape. Mandela wrestled with tension: "Is one justified in neglecting his family on the ground of involvement in larger issues? Is it right for one to condemn one's young children and aging partners to poverty and starvation in the hope of saving the wretched multitudes of this world?" (Mandela, 2018, p. 237). Gandhi wrote: " I am of opinion that all exclusive intimacies are to be avoided.... [H]e who would be friends with God must remain alone, or make the whole world his friend. I may be wrong, but my effort to cultivate an intimate friendship proved a failure" (Gandhi, 1957, p. 19).

One might object: can't the agapic agent love all only very weakly? If her love were faint, maybe she could carve out some space for friends and other particular loves. But consider how we in fact get by in our daily lives despite our knowledge of the dire needs of the many: we largely ignore them. To get by in our days, to avoid being overwhelmed by the suffering of others, we spend much our time not sparing the many in need any thought at all. We restrict our emotional vulnerability—hardening our hearts to most. We might allocate por-

tions of our lives—even large ones like entire careers—to acting on behalf of and maybe feeling with strangers in need, but when it comes time to be present with family or friends, we largely “turn it off,” setting the problems and needs of others to the side for a while. That is, it takes a positively unloving stance towards strangers to maintain the emotional space in our lives for ourselves and loved ones. Whatever (relatively minimal) emotional susceptibility we can have towards humanity as a whole falls well below the kind of emotional openness needed to love.

### 1.6.2 Ideals

What should we make of this crowding out? Though the humanity of others is a reason to love (all of) them, we have strong reasons not to love all, viz., that loving all would prevent us from adequately manifesting the particular loves in our lives, leaving room for little else. I take for granted that if ever we are acting well in our lives, it is when we act out of particular love for family, friends and the like (Keller, 2013). So if manifesting agape rules out our ability to manifest such love, it cannot be required. Moreover, agape would be a mistake, as it would constitute an unacceptable neglect of our loved ones and that which gives our life meaning.<sup>38</sup>

In the passage quoted above, Cottingham goes on to conclude that holding up impartial concern for all as an ideal “make[s] nonsense of ethics” as it “sever[s] the link between ethics and eudaimonia, the good for man or human fulfillment” (1983, p.90). But that is too strong. For those like St. Benedict’s monks who eschewed personal

---

<sup>38</sup>As I argue in my second chapter, loving another involves a commitment to them.

relationships, agape would not be a mistake. (Though neither is it required.) While those who love all foreclose one source of human fulfillment (i.e., particular love), they gain access to another (i.e., plural love). All else being equal, we have reason to pursue either.

But is Cottingham right at least about those for whom all else is not equal? Given that agape would do such violence to our lives and our ordinary loves for friends and family, is it irrelevant? No. We should see love of all as a practical ideal. By an ideal, I mean a trait of character that serves as a standard of evaluation for our actions and attitudes without it necessarily being the case that one should fully manifest it. Insofar as an agent falls short of agape, she is non-ideal—her attitudes and practice are in some way deficient—but she may be justified in so falling short, e.g., by her love for particular individuals. Agents have defeasible reason to manifest the ideal, but do not have all things considered reason to do so because agape crowds out other loves. But the fact that agents have a pro-tanto reason to be agapic (or put another way that the humanity of all provides a pro tanto reason to love them) means that inasmuch as they can, agents should approximate agape consistent with their particular love. Though agape is off limits to those with particular love, we can, and so should, nevertheless be agape-like in our relationships to others.

An ethic that puts the ideal of love for all at its foundation offers a general theory of ethics in terms of agape. So by way of what is required, it would hold that we are required to manifest some minimal approximation of agape. I take it for granted that that which is required has a constitutive link with that which can be demanded by others.<sup>39</sup> That which is supererogatory is that which is more agape-

---

<sup>39</sup>See Darwall (2006, 2013).

like than is required. In general, we should be as agape-like as we can be without running afoul of the reasons not to love all, i.e., without crowding out the room we have for particular love.

The question is then what is involved in approximating agape. What is the minimum approximation demanded? What would it be to exceed that minimum? How agape-like can one be without crowding out room for particular love? A full answer to these questions is, on the view I'm advocating, more or less a full theory of morality. So I will give a sketch of this love-first approach. And here is the promised derivation: from approximations of agape will emerge a familiar-looking deontology. That is, an ethics of love will ground an ethics of respect.

### 1.6.3 Respect as an Approximation of Agape

There is of course much debate among deontologists. But ethical views centered around respect generally have something like the following structural features:

1. Universality: Everyone must be respected.
2. Impartiality: Everyone is equal, and absent special circumstances (e.g., having a special relationship with someone) should be treated equally.
3. Anti-Welfarism: Respect for others requires taking some heed not just of their welfare but of their choices. Respect requires that we honor constraints against promoting another's wellbeing when doing so infringes on their autonomy.
4. Anti-aggregative: To respect the autonomy of all is to see and treat individuals as non-fungible. Respect is committed to the



separateness of persons. In other words, respect requires that we honor certain constraints when making interpersonal comparisons.

5. A distinction between perfect and imperfect duties (or positive and negative rights or something of the sort): An ethic of respect typically involves a distinction between constraints that are always in place, e.g., against harming or lying, and requirements that must be satisfied enough, e.g., an obligation to charity or the development of one's talents.
6. Supererogation: Some actions are such that they are in some way more ideal (or better or we have reason to do them) without their being required.

I'll now argue that an ethical view centered on the approximation of agape has just these features.<sup>40</sup>

It is because agape crowds out room for other parts of a life—especially particular love—that we should not manifest it. That is, the reason that justifies deviation from the ideal is the need for space in our lives that agape would fill. But many elements of agape do not contribute to this crowding out. In particular, one could recognize constraints against promoting an individual's welfare against her will and refuse to reason aggregatively. An attitude towards others that had those two elements of agape would be agape-like without yet causing any problems for particular love. Likewise, you could extend those two elements to all individuals, thereby seeing these constraints as universal. These three elements of agape—universality, anti-welfarism

---

<sup>40</sup>Cf. debates in the Christian tradition about the love commandment somehow explaining the rest of God's commands or agape being at the root of all virtues (Outka, 1972, p. 133-37).

and anti-aggregation—do not pose any obstacle to manifesting particular love. So the justification for deviation from agape, i.e., that agape crowds out particular love, does not justify deviating so far from agape that one does not at least manifest those three features. So, too, with impartiality. While reasoning wholly impartially in the manner of the agapic agent would crowd out particular love, treating all as equal unless there are some special circumstances would not. So one can approximate agape by treating impartiality as the default, from which one can deviate with good reason (e.g., that this person is my friend). Since one can, without dominating a life, reason in such a way that treats all impartially by default and is subject to constraints against harming, lying, etc...or seeing the harms to one as compensated by benefits to another, that extent of approximation of agape is required.<sup>41</sup>

What about agape's promoting of the welfare and autonomy of all? To do as the agapic agent does is untenable; but doing only some of what she does is not. Allocating some fraction of one's time, effort, resources, and affection to serving the welfare and autonomy of all can still leave plenty of room for particular love and other meaningful pursuits. So the approximator of agape is required to do just that. This is all to say that one who acts as required will be following some division between constraints that must always be obeyed (e.g., not to harm) and obligations to aid those in need enough. Exactly how much time, effort, etc..., is enough? This is, of course, a famously difficult question to answer, so I won't do so here. But on this agape-based view, the question will be posed as "how much of an approximation

---

<sup>41</sup>Niebuhr thinks justice is an "approximation" of agape exactly because of its commitment to equality, reflected in the commandment not just to love thy neighbor but to do so "AS THYSELF" (1957, p. 111, 108, his capitalization).

of agape is required?”

One element of agape, though, will not be required: the affective component. Again, I'll help myself to an assumption: that feelings cannot be demanded.<sup>42</sup> While they are subject to standards of evaluation, we cannot demand others feel a certain way, and so (given the constitutive connection between demand and requirement) feelings are not required. That which is required is then an entirely practical approximation of agape. We are not required to feel like the agapic agent does. We are not required to be emotionally susceptible to all. Feeling for others and going beyond whatever minimum standard of aid is required are both more like agape and so more ideal. That is, they are supererogatory.

An ethic of agape, then, has all six of the features identified above. An ethical view that takes agape as an ideal to be approximated just is deontology. But instead of taking respect for all as a first principle or respect as, in the first place, called for by the humanity of others, it sees the need for respect as derivative.

## 1.7 Love First

Recall the question with which this paper began: how should we respond to the humanity of others? I've argued that fundamentally, the answer is love. The humanity of others calls on us as agents to

---

<sup>42</sup>“[L]ove as an inclination cannot be commanded, but beneficence from duty...is practical and not pathological love, which lies in the will and not in the propensity of feeling, in principles of action and not in melting sympathy; and it alone can be commanded” (1998a, AK 4:440, *emph. original*). Thus, for Kant, a duty to love is “an absurdity” (2017, AK 6:401). But following God's command to love (practically) just is, for Kant, a command to follow the moral law (1998b, AK 6:124). The structure of his view, then, is the reverse of mine: for Kant, we begin with the moral law (of respect), which then explains the content of the injunction of practical love.

respond with love, and when we cannot, to fall back on an approximation thereof: respect. The familiar structures of deontology are given a novel ground in love. The argument for all that has three parts. To present them out of order, one part is the articulation of agape in light of what seemed like conceptual or practical impossibilities. Another is the construction of a view that takes agape as its central ideal, giving an account of what is required and what one should do in terms of that ideal. Together these amount to getting a love-first approach to ethics in view. The third part—the argument in favor of that approach—came at the beginning: my claims about love’s outward pressure (§1.3). I argued that when we thought about ordinary cases of love, we saw a pressure towards loving all. It is this pressure that serves as the basis for claiming that love is an ideal that we should approximate inasmuch as we can.

With the whole picture in view, it is worth returning to that initial motivation towards agape. I claimed that what is good about love is its expansion of the heart and connection with others, and so what’s good about love admits of expansion and gets better the more one loves. And I argued that the rationality and moreover ideal character of unconditional love shows that the humanity of each person gives us reason to love each human being and so all human beings. But now that we’ve seen where this pressure leads, one might return to the objection voiced at the end of §1.3. As Velleman puts it, “respect and love [are] the required minimum and optional maximum responses to one and the same value,” that of another’s humanity (1999, p. 366). But why see the requirement of respect as derived from agape and not as matter of first principle? The alternative is a view on which humanity requires respect and makes love rational without giving any

(insistent) reason to love. On this view there is no pressure towards loving all nor any explanatory route from love of all to the need to respect all. Given that acceptance of rational pressure towards love paves the way for the rest of the paper, asking why we should take love as an insistent reason for love as opposed to a noninsistent reason that merely rationalizes love is a direct challenge to the whole explanatory strategy I've offered; asking whether humanity is an insistent reason for love is asking why we should go in for the love first view I've articulated.

I am now in a position to make good on the promises I made in §1.3. Where Velleman and others dismissed or had arguments against the possibility of loving all, we've now seen how to make sense of the claim that the humanity of others gives us a(n insistent) reason to love all while still maintaining that ultimately, love of all is at best optional and for many of us a mistake.

Second, we can now see that a love-first view offers an explanatory unity that is otherwise elusive. "Love, like respect, is the heart's response to the realization that it is not alone," (Velleman, 1999, p. 366). But why are love and respect so intertwined? Why are they the minimum and maximum responses to the same thing? What does love have to do with respect? To see each as a response to the same thing and as arrayed on a spectrum from maximum allowed to minimum required is to invite the question of their relation. And the love-first view has an answer: respect is love's approximation. Seeing the humanity of the other as pressuring us towards love then enables us to explain that spectrum in two ways. First, it takes only love as a starting point and derives respect as love's approximation. We have only one normatively fundamental attitude and define the

other in terms of it. Second, it explains the optionality of the one and requiredness of the other from the same starting point, viz. that the humanity of each is a reason to love him. In so doing, it brings together in a common source our obligations towards the many and the special attitude we have for only a few—the fundamental equality of all and the partiality with which we are permitted and obligated treat a select few can both be explained in terms of love.<sup>43</sup>

A view that takes respect as underived cannot make similar moves. In general, it is hard to see how one could derive an optional maximum from a required minimum—that would be like trying to understand beauty in terms of mediocrity. More specifically, it is hard to see how we could derive or analyze love in terms of respect. If analyzing knowledge as true belief plus some elusive *x* is hard, the prospect of defining love as respect plus some elusive *y* looks even worse. The respect-first theorist will have to say that the humanity of another's justification of love and requirement of respect are two independent normative upshots of the same fact. That is suspicious—it calls out for some unifying explanation. And while a failure to explain a salient connection is not a decisive objection to a respect-first view on its own, it is decisive when its rival does explain it.

Third, accepting a pressure towards love as opposed to seeing it as merely optional better captures the phenomenology of living a moral life. Even if we act at our best, we still, in some hard-to-identify sense, seem to fail. However much we do for others, we could do more. And though we may be confident that we have done enough—even that if we were to do more we'd have to sacrifice that which should not be sacrificed, like our ability to manifest love for friends and family—

---

<sup>43</sup>See chapter 2.

there remains a feeling of our failing to live up to something. Acting as well as we can is still acting non-ideally. If only we could respond more fully to the humanity of others..., but we can't, at least not without thereby making some other mistake. Failure, then, seems unavoidable. The nature of the question posed by the existence of other people—how should we react to the humanity of others—does not have a wholly satisfying answer.

Seeing agape as an ideal captures this hard-to-identify sense of inevitable failure. It is not a case of doing what we should not do or reasoning incorrectly. But it is nevertheless a deviation from an ideal. At the same time, it is an ideal out of reach for most of us, one we cannot but fall short of absent a total abandonment of particular love—which would leave our life deficient in another respect, not to mention constitute an abandonment of those to whom bear love's commitments. Living virtuously is a matter of living in this shadow of agape, doing the best we can to reach towards it—to fully respond to the humanity of all—and always falling short. And this is at once both depressing and liberating. Responding well to the humanity of others is a matter of living up, as best we can, to the almost unreachable standard of perfect love for all. But though impossible, it is a standard whose intelligibility and normativity can be found in the part of our lives where we can be most sure we are getting something right—in the love we bear for our family, our partners, and our friends.





## Chapter 2

# The Importance of Being Constant

### 2.1 Partiality

Why spend money on your child when others are in need? Why visit your friend in the hospital when that time could be spent helping the indigent? Suppose a man (“Bernard”) is standing on a pier with his wife (“Patricia”) and a stranger drowning on opposite sides. Bernard can save only one of them, and he saves his wife (adapted from Williams 1981). In the years that come, Patricia and Bernard remain at the center of each other’s lives; while not selfish, much of their time, money and effort is spent on one another. They often do this without sparing a thought for those strangers they might have otherwise benefited. Spending the money, visiting the friend, Bernard’s saving Patricia, their building their lives around one another—these are instances of partiality. A theory of partiality must offer an account of both the moral psychology and justification of partiality; but as we will see, these two desiderata pull in seemingly irreconcilable directions. I argue that this dilemma can be resolved by appealing to the significance, rational and psychological, of history. Virtuous

partiality involves the adoption of a principle of decision-making over time; we are permitted and sometimes required to be partial in large part because of our history of decision-making.

At a first pass, partiality is an agent's bias in favor of some select few. For an agent to be partial is to have the interests and preferences of the few in some way count for more in the her deliberation—one way or another, it is the interests of the few on which she acts and in response to which she forms attitudes of care and attention. Any talk of partiality makes salient some kind of conflict between partiality, on the one hand, and impartiality, on the other. It is a conflict with which philosophers are familiar, and on which there is more to be said, though not here.<sup>1</sup> Though these issues will loom large in the background, we will not consider them directly. I'll take for granted that at least some partiality is virtuous (or "rational" or "the thing to do"—whatever you like) even in the face of the urgent needs of strangers. It is the task of this paper to offer a theory of partiality itself.

What does such a theory need to explain? Two things: the moral psychology and the normative structure of partiality. That is, what is it to be moved to act in the way that is distinctive of (virtuous) partial behavior? And what makes it the case that an agent may or must act in a partial manner ("act partially" for short). When Bernard saves Patricia, why does he save her? And what makes his action permissible, and moreover, required? In the hard cases where it is not clear whether we ought to be partial or not, what moves us toward partiality, and what facts speak (normatively) on behalf of acting partially?

---

<sup>1</sup> See the first chapter.

Our theory must answer to these explanatory demands—the normative and the motivational—and it must do so in a coherent and unified way. The trouble, illustrated by extant theories of partiality, is that these two desiderata pull us in opposite directions. On the one hand, relationships seem like an essential part of the normative story. Patricia and the stranger on opposite sides of the pier are (in some sense) equal; it is not as though Patricia’s life is inherently more important. The reason Bernard should save her instead of the stranger surely has everything to do with the fact that Bernard loves Patricia, that she is his wife, etc.... It is these relational facts that make a normative difference. On the other hand, the thought “Patricia is my wife” is one thought too many. As I discuss in §3.1, there is a certain kind of direct, unmediated loving action that precludes a thought of the relationship, and such action is virtuous. Assuming a connection between a virtuous agent’s motivating reasons and justifying reasons we have a dilemma: Bernard’s relationship to Patricia seems necessary to justify Bernard’s action, even though when he acts virtuously, Bernard is not motivated by any relational fact. There is a tension between the normative and motivational desiderata, with the former saying the relationship must play a role in the ethics of the situation and the latter suggesting it does not play a role in the agent’s psychology.

Extant views tend to be partisans of one desideratum at the expense of the other. So-called “relationships views”<sup>2</sup> claim that relationships are the reasons that ground partial permissions and obligations: Bernard should save Patricia because she’s his wife and he loves her. But by appealing to relationships, they yield false implications

---

<sup>2</sup>I take the term from Keller, whose taxonomy in his 2013 is insightful and illuminating.

in moral psychology, namely, that beliefs about relationships play an essential role in (virtuous) partiality. In response, “individuals views” try to do the normative and ethical work without an essential appeal to the significance of relationships (e.g., that she’s his wife or that he loves her). Instead, they claim that non-relational facts (e.g., that Patricia is drowning) justify partial action on their own. But in so doing, individuals views get the moral psychology right and the ethics (mostly) wrong, unable as they are to draw adequate distinctions between strangers and those to whom an agent bears special relationships. We have special permissions and obligations to be partial, and individuals views cannot adequately account for either.

Faced with this antinomy, I’ll defend a novel view. I’ll argue that the debate has rested on a mistake: an assumption that the rationality of partiality is a wholly synchronic matter. When an agent acts partially, others assume that it must be something about her psychology at that time in virtue of which her partiality is rational. But partiality is not a phenomenon that can be understood without recognizing the centrality of agents’ histories of being partial. Permissible cases of partiality involve seeing and treating someone as especially important over time; it involves adopting a principle of practical reasoning that gives them special weight in one’s reasoning. I’ll argue that we are under an essentially diachronic, defeasible requirement of rationality to be constant with respect to such practical principles. A view whose central thesis is a requirement to be constant across time is able to square the circle, preserving what is right about the moral psychology of individuals views while nevertheless having the richness to draw necessary moral distinctions between strangers and the special relationships that generate special permissions and obligations of

partiality.

I'll begin by considering relationship and individuals views in turn, developing them and the objections they face in greater detail both for the sake of setting up my own proposal and in the hopes of mapping the landscape of a still under-appreciated debate. I'll then defend my appeal to history and defense of a rational requirement of constancy. To bolster the case for constancy, I'll show that adoption of the view can shed light not only on the particular case of love but more general puzzles regarding decision-making under conditions of parity.

## 2.2 Taking Love Seriously

To investigate partiality, we will look towards love. I'll assume that an agent who acts (virtuously) out of love is motivated by that which in fact justifies (or requires) her action. This connection between the moral psychology and ethics of loving action has served as methodological common ground in discussions of partiality, and virtually all parties to the debate considered in this paper have accepted it.<sup>3</sup> Why? Keller puts it in terms of special relationships but his point is stronger if we read "special" as "loving":

The times when we act well towards those with whom we share [loving] relationships—when we respond immediately to the needs of our parents, make a sacrifice for the sake of a friend, or do something that brings joy to

---

<sup>3</sup>The connection was given canonical voice in Williams 1981 and Stocker 1976. The contemporary authors to whom this paper objects all rely on or at least endorse a connection between an agent's treating or seeing a relationship as a source of reasons (or as reason) and its actually being a source of reasons (or a reason). Scheffler 1997, 2010; Kolodny 2003, 2010a,b; Jeske 2008; Keller 2013; Pettit 1997; Setiya 2014; Velleman 1999.

our children—are some of the times when we seem to see most clearly what matters, when we seem to be in closest contact with the really important things in life. It would be both implausible and depressing to suggest that when we act well within [loving] relationships we systematically misperceive our reasons. It is difficult to see what sort of philosophical argument could convince you that when you act out of your special concern for your friends, your partner, or your family, the things that really provide you with reasons are different from the things that seem to provide you with reasons. (2013, p. 27)

Sometimes it is hard to know whether we are doing the right thing, but caring for those we love is not one of those times. If ever we are getting things right, it's then. The facts that a virtuous, loving agent takes to justify her actions really are sufficient to do so; the facts that strike a virtuous, loving agent as inappropriate considerations on which to base her action really are so. So if we can show that a consideration seems necessary, unnecessary, appropriate or inappropriate to an agent acting out of love and doing so virtuously, we've learned something about the ethics of partiality.

### 2.3 Love and Relationships Views

Relationships views place, well, relationships at the center of a theory of partiality. They thereby get the ethics of partiality right, explaining just when we are permitted or required to be partial: when we stand in (certain) special relationships.<sup>4</sup> It is to our friends, our partners,

---

<sup>4</sup> See Kolodny 2003, 2010a,b; Scheffler 1997, 2010; Jeske 2008.

our children, our loved ones, our colleagues, etc..., that we should be partial, and that is so, they say, because of the relationships we bear to them.

Let's return to our original case in which Bernard can save his wife or a stranger. Relationships views will claim that it is because Bernard's relationship to his wife is so valuable (or provides such strong reasons) that he is required to save her over the stranger, and would even (say) be permitted to save her over the lives of a small handful of strangers. It is because the relationship is not valuable enough that Bernard would be forbidden from saving Patricia at the expense of many thousands of lives. And it is something about the difference between (say) being a partner with and being a co-citizen with that explains why the former justifies and requires far more partiality than the latter.

With respect to the moral psychology, relationships views claim that a thought about the relevant relationship ought to feature in an agent's practical deliberation as the content of some motivating attitude. The relationship might make an appearance in a belief that the agent shares a special relationship with the patient, where that belief serves as the basis (or part of the basis) of the agent's decision to act partially. The belief might be more complex, involving some explicit valuation of the relationship as a source of reasons or as good in and of itself. So, too, the thought might involve more than a belief, perhaps a desire or complex set of mental states. The point is that on a relationships view, the relationship itself figures as some kind of basis for the virtuous agent's partial behavior.

When it comes to the ethics of partiality, something about a relationship-based approach to partiality—and love in particular—

almost has to be right. We seem to be justified or required in acting partially just when we bear some kind of special relationship to the beneficiary of our actions, so it is natural to think that the relationship itself plays a normative role. When faced with a choice to save a loved one or a stranger we are required to save our loved one, and the fact that we share this loving relationship seems to be the obvious fact that tips the balance of normative reasons decisively in favor of the beloved.

### 2.3.1 Relationships views disallow directly motivated action

Despite these strong presumptions in favor of some kind of relationship-based view of partiality, thoughts about the relationship do not seem necessary when acting out of love. As Parfit observed, “It’s odd that Williams gives, as the thought that the person’s wife might hope he was having, that he is saving her because she is his wife. She might have hoped that he [would save] her because she was Mary, or Jane, or whatever. That she is his wife seems one thought too many” (Parfit as quoted in Murphy 2000, p. 140, n36). Frankfurt puts the point even more strongly, “I cannot help wondering why the man should have even the one thought that it’s his wife. Are we supposed to imagine that at first he didn’t recognize her? Or [that] he didn’t remember that they were married, and had to remind himself of that? It seems to me that the strictly correct number of thoughts for this man is zero” (2004). While Frankfurt’s claim that a husband of good character should be literally thoughtless while saving his wife is somewhat overstated, he and Parfit are on to something important. It is not just Williams’s intended targets that are subject to worries about motiva-



tion and partiality. When Bernard jumps into the sea to save Patricia, it seems that he need only be motivated by a singular thought about Patricia. Patricia could hope that in his love, Bernard is moved to save her solely by thoughts about her, and without thinking about about how she relates to him or about the relationship itself.

Pettit writes:

It is doubtful whether I could claim to be properly a lover, if it was my recognition of the fact of loving her... which explained my action: if all that needed to be said in explaining how I behaved was that I saw I loved her or saw that I bore a relation to her which, as it happens, means that I loved her.... To act out of love, as we might put it, is to be moved by love and not by the recognition of love.

(1997, pp. 155-56)

He puts the point too strongly. There is nothing “doubtful” about the agents’ love in these cases; indeed, real, loving relationships are filled with actions in which a thought of the relationship prompts us to do more for those we love. A couple facing difficulties in their partnership might find that they are motivated to work through them so as to “save the marriage.” No worry about having “one thought too many” seems to apply to agents who thoughtfully value and tend to their relationships in these ways. Moreover, many of the difficult things we do for our loved ones involve being motivated by the thought that these relationships require us so to act. Nor do such actions fail to express love.

But there is a further kind of acting out of love—the kind of action out of love through which (I submit) we most purely express our love and by which we most clearly feel loved—what Pettit takes to

be love simpliciter—wherein the beloved’s needs, interests and preferences alone suffice as motivating reasons for the agent. It is a kind of love that is expressed when an agent is wholly directed at and moved by her beloved, rather than by the fact of her love or relationship. When Parfit writes that the wife in Williams’s case “might have hoped that he [would save] her because she was Mary, or Jane, or whatever,” he is giving voice to something like that thought. It is the hope that you, as the student or partner in the previous examples, might have: that as soon as your parent or loved one recognized your need it would be sufficient to motivate them to act, without requiring recourse to a thought about the relationship they bear to you—that you are enough. If Bernard saves his wife without being moved by their relationship, he are not missing a thought, the absence of which renders the action in some way deficient. Relationships views falsely claim that one who so acts without any thought of the relationship is acting on insufficient reason.

The objection can be put more or less stridently—we can think (as I do) with Parfit that such directly loving action is not just permissible but preferable, and complain that relationships views disallow this form of virtuous, loving action. Or more modestly, we can claim that it is at least permissible to so act; directly loving action is not action performed in some error. And while there is more to be said on the matter,<sup>5</sup> either version of this “one thought too many” objection poses a serious problem with the way in which relationships

---

<sup>5</sup> For an excellent defense of a relationships view in light of these kinds of objections, see Kolodny 2003, §6. For a decisive response to Kolodny and further, carefully argued objections to relationships views, see Keller 2013, and Setiya 2014, whose defense of an individuals view rests on the rationality of unconditional love, of love at a distance, and of love that survives of loss of belief in the relationship, e.g., through amnesia.

views incorporate relationships into the moral psychology. We must find some other means of incorporating what is right about relationships views, viz., that the rationality of partiality turns somehow on the significance of special relationships, without taking relationships themselves to be the basis of partial action.

Note in closing that the problem is not that acting because of the relationship is acting “too reflectively,” as though the problem were that such love is not sufficiently spontaneous.<sup>6</sup> The claim is instead that there is a kind of direct, unmediated loving action in which an agent spares no thought for the relationship itself as a condition, basis or reason for her action. Such action is virtuous—at least permissible and perhaps preferable—and relationships views make it out to be some kind of mistake.

## 2.4 Individuals Views

Inspired by thoughts like these, defenders of “individuals views” offer theories of partiality that do without an essential appeal to relationships.<sup>7</sup> They are united by the core idea that to act out of love does not necessarily involve being moved by thoughts about one’s relationship to the beloved. Loving agents can be moved solely by thoughts about the individual herself. Accordingly, they try to make sense of the idea that non-relational facts, e.g., that Patricia was in need, suffice to justify partial action, e.g., Bernard’s saving her over the stranger. Relational facts might show up somewhere in the moral

---

<sup>6</sup> As it is construed by Jeske 2008, p. 130-31.

<sup>7</sup> See Velleman 1999; Setiya 2014; Pettit 1997. Keller would say individuals views are those on which the reasons on which we act make no essential reference to relationships; he would therefore count his 2013 as an individuals view, though, as we’ll see below, I doubt it should be counted as such.

psychology; for instance, Bernard's love of Patricia might be a disposition of character to take non-relational facts involving Patricia's wellbeing as especially weighty or important in deciding what to do. But the relationships do not bear the justificatory or normative weight in the theory of partiality.

Individuals' views get the moral psychology largely right, but at a steep price. In short, they do not even seem to be in the business of explaining the ethics of partiality. Without the appeal to relational facts, the theories cannot draw the normative distinctions between strangers and those to whom we may and must be partial, i.e., those to whom we bear certain special relations.

#### 2.4.1 Getting the Moral Psychology Right

Bernard need only be moved by the thought "Patricia needs help." Individuals' views start with the idea that such non-relational thoughts alone suffice to motivate Bernard, and therefore to justify his actions. Does that mean relationships are entirely absent from the picture? Even just considering the motivations of Bernard, that would be implausible. After all, if we asked Bernard why he saved Patricia and not the stranger, he might naturally respond, "She's my wife"; if you asked me (a third party) why Bernard saved Patricia, I might likewise respond, "She is his wife." Likewise, if Bernard explained his action by saying that Patricia was in need, it could be natural to challenge him, saying, "but so, too, was the stranger!" The obvious rejoinder on Bernard's behalf is that Patricia is his wife.

Does the fact that we explain action with such appeals to the relationship tell against individuals' views' motivational claims? No. There is a difference between a reason why someone acts and the rea-

son why something is a reason why. That's a mouthful, so consider an analogy: We might explain a courageous or generous person's action by saying citing her virtue of character (why did he save that child? because he's courageous); it does not follow that she acted on the basis of her virtue. Her courage is not the reason why she acted; instead, citing the virtue is a way of saying why the facts that moved her (e.g., that child was in danger) did move her. Likewise, citing the relationship cites the practical principle, which, like a virtue, makes clear why the reason at play (Patricia's need) moved the agent: because an agent with a practical principle of loving Patricia will be moved by individualistic facts regarding her.<sup>8</sup>

An individuals theorist, then, can claim that the relational element appears in Bernard's moral psychology in just the same way a virtue does. To love Patricia is like being generous. A generous person is not moved by the fact that she or some action of hers would be generous; she is moved by the needs of others. So, too, to act in genuine love for Patricia is not to be moved by the fact that one is or one's action would be loving; it is to be moved by her needs. To act because of the love (or relationship itself) is like acting because doing so would be generous: it is continence, not virtue, displaying "a commitment to love rather than a lover's commitment," or a commitment to a relationship rather than the commitment of the relationship (Pettit, 1997, p. 156).

What exact role do virtues or love for another play? One might think that in general, virtues like courage and generosity are not merely dispositions to be moved by certain facts qua reasons. Such a view might instead require that the disposition be essentially ratio-

---

<sup>8</sup>cf. Stocker 1981.

nal,<sup>9</sup> or require that the agent adopt a maxim and issue some kind of self-legislation,<sup>10</sup> decide to treat something as a reason,<sup>11</sup> or deciding on some relevant end.<sup>12</sup> Thus, it might be that to love a friend involves not a disposition, but in some way seeing yourself as a friend, seeing your friends' needs under a distinctive normative guise (e.g., as reasons for you), deciding to take your friends needs as reasons, or setting your friends' wellbeing as an end.<sup>13</sup>

The differences between these various conceptions of action and practical reasoning are important, but not for our project. So I'll use the term "practical principle" or "principle of practical reasoning" as neutral between any such views.<sup>14</sup> To some, that's just a disposition; to others its much more. So long as the notion at play does not, in the case of virtue, make something like courage a premise of the agent's practical reasoning, it can play the kind of role that individuals theorists would have it play for love and partiality: an important aspect of the moral psychology that does not feature in agent's first-order reasoning.

## 2.4.2 Getting the Ethics Wrong

Motivationally, individuals views are on good footing. But normatively, they are in trouble.<sup>15</sup> Partiality involves special permissions

---

<sup>9</sup>As in Wedgwood 2006 or Arpaly and Schroeder 2015

<sup>10</sup>As in ? and the work of Kantians like Korsgaard 2009.

<sup>11</sup>As in Bratman 1999b.

<sup>12</sup>As in Bratman 1999a.

<sup>13</sup> cf. Korsgaard 1996.

<sup>14</sup> I'm partial to the kind of dispositional accounts defended in Wedgwood 2006; Arpaly and Schroeder 2013, 2015; Setiya 2011 and Setiya 2010, who, despite important disagreements, all adopt some kind of dispositional approach.

<sup>15</sup> It is no accident that Velleman 1999, which defends an individuals view about the emotion of love, does not even attempt to extend that account into an account of what reasons we have to act. On his view, the mere fact that someone is a person is reason to love her; that one bears a special relationship is besides

and obligations to a select few. A theory of partiality must explain the difference between strangers and Patricia for Bernard. He is permitted to do more for her than for strangers; and moreover he is obligated to do for her that which he need not do for strangers. Let's take each of these issues in turn.

First the permissibility problem: Bernard may be moved more strongly by facts about Patricia than about a stranger. Imagine a pier scenario with Patricia on one side and some number,  $n$ , of savable strangers on the other. There is some value of  $n$  such that Bernard, who loves Patricia, may still save Patricia, but a third party to whom all are strangers would be acting impermissibly in saving the one over the many. You might think the number is simply two; but individuals theorists have made much of the parity of reasons.<sup>16</sup> It may be that the reasons to save one life are on a par with the reasons to save two, i.e., when it comes to the value of lives, reasons don't just cleanly add up such that two lives provides greater reason than one.<sup>17</sup> Nevertheless, at some point, it would be impermissible for an agent to save one stranger over many where Bernard would still be permitted to save Patricia. That is just what it is to have special permissions with respect to our loved ones—we are permitted to favor them at the expense of others in a way we cannot with respect to strangers.

Individuals views cannot adequately explain this difference be-

---

the point. But partial action, he claims, is justified and required by the fact of special relationships and the "mutual commitments and dependencies" thereof (1999, 108). But as others have noted (cf. Kolodny 2003; Setiya 2014) in divorcing reasons for love from reasons for loving action, Velleman's view suffers a kind of disjunctive schizophrenia.

<sup>16</sup> See especially Setiya 2014.

<sup>17</sup> To say that reasons are on a par as opposed to equal means that parity between the reasons is not broken by small or even medium-sized adjustments to the strength of the reasons. For more on parity, see Chang 2002, the introduction to Chang's 1997, and Hare 2010.

tween

strangers and loved ones. After all, take the value of  $n$  such that Bernard may save Patricia but a stranger in his situation may not. In keeping with their motivational arguments, we should allow that Bernard's saving Patricia over  $n$  strangers can be a directly motivated, loving action: he is moved to save her solely by non-relational facts about her needs, i.e., facts solely about Patricia. He is faced by two sets of reasons: those involving Patricia, which tell in favor of jumping to her side, and those involving the  $n$  strangers, which tell in favor of jumping to the other side; he permissibly acts in accordance with the former. But were it not Bernard but another stranger standing on the pier, facing the decision to save 1 stranger (Patricia) or  $n$  strangers, he would be faced with the same set of reasons. If Patricia's needs alone justify Bernard's saving her, so, too, they should justify a stranger's saving her. But they don't. And the obvious explanation of Bernard's special permission—that Bernard loves Patricia or that she is his wife—is exactly the fact that individuals views are supposed to do without in explaining what moves Bernard.

Individuals theorists are not unaware of this problem. One possible response is to allow that while some partial action is directly motivated, not all of it is (Setiya, 2014). Sometimes, virtuous agents are motivated by relational facts. So while Bernard might be motivated by non-relational thoughts about Patricia in a 1 vs. 1 case, he must base his action in part on the relationship in a 1 vs.  $n$  case. The difference between the permissions we have to favor strangers and to favor loved ones can be explained by an appeal to the relationship as a supplemental reason.

But this response will either be insufficient or open the view up



to the very criticisms against relationships views that it was trying to avoid. After all, most of the partiality we exhibit is a lot more like a 1 vs.  $n$  case than a 1 vs. 1 case. We might give a close friend an expensive gift when that money could have paid for the deworming of hundreds of children, greatly increasing their chance of finishing secondary education. Such actions are permissible, and when agents exhibit the direct motivation distinctive of love to do so, they are not acting on insufficient reason. But we may not so favor the one over the many when it comes only to strangers. Choosing to give an expensive gift to a single stranger instead of spending the money on deworming, supporting a museum, or something of the sort would be a objectionable singling out of one among many equally worthy recipients of charity.

Individuals theorists are faced with an dilemma. They can say that agents who favor 1 over  $n$  in a directly motivated manner are acting on insufficient reason. The virtuous gift-giver, they can say, must reason a la relationships theories. But this has the unacceptable result that any such action that is directly motivated by thoughts of the beloved alone is done for insufficient reason.<sup>18</sup> Or, they are forced to say that a case like that of giving an expensive gift to one over the

---

<sup>18</sup>Keller's proposal that in all of these cases, the relationship plays the role not of a reason but a modifier falls prey to this worry (2013). Grant reasons and modifiers are normatively different. On this view, the relationship is nevertheless a fact that must be recognized by the agent to justify partial actions, even if the motivating thought's role is as a modifier and not a reason. Bernard would still need to have his relationship with Patricia form (part of) the basis of his action. Were he to lack this thought, the attitudes that motivate him (a belief in Patricia's need) would, on this view, be insufficient to justify or obligate partial behavior—saving Patricia instead of strangers. But this clearly won't do. The objection wasn't that Bernard had "one reason-thought too many," it was that he had one thought—of any kind!—too many. Keller's faces the same objection as relationships views: the fact of the relationship is what makes the normative difference, and an agent is reasoning well only insofar as she bases her action on the facts that justify them, which include the fact of the relationship.

many isn't irrational for anyone to do. It is behavior that any stranger is justified in doing for any stranger, disparity in outcomes notwithstanding; much of what we had thought of as special permissions are permission that anyone has with respect to anyone else. That is to say, this is a view on which the numbers really don't count in ethics.<sup>19</sup>

The obligations problem is much the same: without an appeal to relationships, individuals theorists cannot explain the special ways in which we are obligated to loved ones and not to strangers. Perhaps in light of the challenge posed by anti-partialist arguments like Singer (1972) and Unger (1996), many discussions of partiality have aimed to show why it is that we are permitted to favor our loved ones and other special relations. One could be forgiven for thinking that the defender of partiality need only establish that we have a prerogative to favor some at the expense of what is impartially best. But if we reflect not on some challenge directed at partiality but our ordinary practices of love, it is clear that many of the norms of partiality are obliging. We owe our friends, parents, siblings, children, etc...all sorts of things. If in a 1 vs. 1 case, Bernard's pier were flanked by two strangers, he would be permitted to save either; but it isn't. Given that one of the drowning people is Patricia, he is not merely permitted to save her, he must.

There are two important explananda: first, Bernard is not merely permitted but obliged to save Patricia; second, Bernard's motivational state will reflect that obligatory character inasmuch as he's acting virtuously. That is, he would not see the option to save the stranger as a live or open one. If asked to explain himself after the fact, he might naturally say "I had no choice" or "saving that stranger wasn't

---

<sup>19</sup> A conclusion accepted by Setiya.

an option—I wish it were and feel enormous regret over his death, but Patricia was in danger.” If reasoning virtuously, Bernard would feel obliged, experiencing Patricia’s need as a reason that does not merely justify saving her over the stranger but requires it. Nor is this phenomenology unique to pier-like scenarios. We often (rightly) feel the needs of our loved ones not merely as reasons on a par with others; they feel like reasons to which we must respond; we are obliged to act on them, and they present themselves as obliging. Not to feel this is to miss something central to one’s practical situation. And we, along with Bernard, feel this obligatory force when we are moved solely by non-relational facts.

But individuals views cannot explain the obligatory character of Bernard’s action. Nothing about the non-relational facts, viz., that Patricia is in need and that the stranger is in need, could explain why Patricia’s needs oblige where the stranger’s do not. The individuals theorist simply has no resources on which to draw to explain why these two reasons that seem to be on a par would have such different effects on a virtuously reasoning Bernard. As above, we can imagine the relationship theorist appealing to the relationship as a supplemental reason. They could propose that while the non-relational facts are sufficient to permit saving Patricia over the stranger, the relational fact then steps in as an additional reason that then obligates Bernard.<sup>20</sup> But this runs afoul of the second explanandum. When Bernard acts in a directly motivated way—just on the basis of Patricia’s need and without being motivated by thought of the relationship—he will feel that reason as obligatory. In other words the only way an individuals theorist seems able to explain the fact that Bernard is obliged lands

---

<sup>20</sup> The following remarks apply just as much to a view like Keller’s on which the relationship plays the role of a modifier, not of a reason (2013). See n. 18.

the view in just those troubled waters it sought to avoid by turning away from relationships in the first place.

To explain both special permissions and special obligations, we need to appeal to relationships. But it seems impossible to do so without running headlong into the motivational problems with the straightforward appeal employed by relationships theorists. The relationship must make a normative difference while not showing up in the content of a motivating attitude. But what could make a difference to what it is rational to do other than the attitudes by which we are motivated?

We can consider one last attempt: might the fact that Bernard's practical principle reflects the relational fact of his love for Patricia do the work? The fact that Bernard loves Patricia (a relational fact) is encoded, as it were, in Bernard's moral psychology as a the practical principle to take special note of her in his practical deliberation. Could the fact that he has such a principle (and so loves her—a relational fact) explain why he is permitted and obliged to save Patricia? Pettit argues that the fact that Bernard reasons in the manner constitutive of loving Patricia explain why the reasons having to do with Patricia have a special weight, both normatively and motivationally (1997). But it is hard to see how this view does not run into problems of circularity. We are seeking to explain why Bernard may, and sometimes must, be more strongly motivated by facts about Patricia; the fact that he has a practical principle according to which her needs and preferences are especially weighty cannot explain why he may and moreover must count her needs and preferences as especially weighty.

## 2.5 An Appeal to History

Here things seem to be at an impasse. But I will argue that we can identify and reject an assumption all parties to the debate have taken for granted: that rationality is an essentially synchronic matter, i.e., that the only things that could make a difference to what it is rational to do at some moment are facts about an agent's psychology at that moment. Though we've run out of plausible strategies that appeal to an agent's psychology at a time, but we can look elsewhere: to an agent's past.

To love another is not to jump off a pier to save them once; life serves up piers of various sorts daily, and we must constantly choose whether to favor our loved ones or others. To love involves choosing our loved ones repeatedly. It is not a momentary attitude; to love someone is to see and treat her as significant over time. The key to understanding partiality is to appreciate the diachronic requirements we stand in, among them, a defeasible requirement to be constant in our love over time.

Consider a case of irrational impartiality. Imagine that after years of loving Patricia, Bernard is faced with a choice between Patricia's life and a stranger's and he gives Patricia's no extra weight, instead just flipping a coin and saving the stranger's life. After years of reasoning in a manner especially attentive to Patricia, he just suddenly stops. Something has gone wrong; what? We don't want to appeal to Bernard's failure to countenance his relationship with Patricia as a reason—that will just land us in familiarly troubled water. What is wrong seems to be his sudden, seemingly crazy shift in the way in which he approached the situation. Those who know him would wonder why he would just stop caring about Patricia? That is, a nat-

ural reaction to this story is to demand an explanation—a reason—for his shift in the way he perceived his practical situation and reasoned about it. Why did he stop paying special attention to Patricia? Why did she all the sudden stop counting for more? Perhaps there could be good reasons for this, but suppose in this case there aren't; Bernard's behavior then seems irrational.

Nor does that sense of something going wrong attend only to a case in which he saves the stranger's life. Suppose he flips the coin and saves Patricia. Still, we would wonder at why he suddenly stopped treating Patricia partially. What seems off about Bernard is not just his choice of whom to save but the way he reasoned about the situation, i.e., his principle of practical reasoning.

Of course, had Bernard lacked any history with Patricia, had he not for years been reasoning in a way that singled her out as among the special few to whom he was partial, Bernard's coin flip would be unobjectionable or maybe even required. So it is not the principle, per se, that is objectionable. It is, I suggest, change—from having a principle that is partial to Patricia to no longer seeing her that way. Imagine Bernard had a short history with Patricia—they just met a few weeks ago. His behavior is less bizarre than if he had spent years with her. All that is to say that it is history that makes a rational difference. Relationship theorists claimed the relationship is what motivates and justifies partiality; but it is rather the moral psychological embodiment of the relationship—the agent's principle of relating to the beloved as special—that makes a difference. It is not a present belief about a historically extended entity, i.e., the relationship; rather, the difference maker is the agent's history of so relating.

The core of the proposal is this: our principles of practical reason-

ing constitutive of loving partiality enjoy a kind of rational inertia. Those that we have rationally adopted require some reason to abandon. And the greater the history of so reasoning, the greater the required reason. To just give up a practical principle is to manifest the vice of inconstancy; it is to violate a diachronic requirement to be constant in our manner of reasoning. The greater our history with a principle, the greater the reason required to justify its abandonment. Our past exerts a rational influence over our present. Of course, the past is not binding. We face some rational pressure to be constant, but that pressure can be overridden. It could be that loving Patricia becomes too difficult and that Bernard should give up his love for her. But that would only be rational because the requirement to be constant was overridden by the difficulties of maintaining his love.

More needs to be said about the strength and shape of this requirement (see §2.7.1), but with the outlines in place, we can see that a diachronic requirement to be constant can solve all the problems we have faced. Recall that relationships views get the ethics of partiality exactly right; once we find a way for relationships to make a difference, we solve the problems about permissibility and obligation that individuals views could not. We have effectively gotten relationships into the picture without locating them as a reason or modifier or anything else that would see them featured as the content of a motivating attitude.<sup>21</sup> A principle of practical reasoning of treating one's friend,

---

<sup>21</sup> In correspondence, Keller helpfully objected that on my view, action out of love will be too self-centered, likening my view to what he criticized as "projects views" of partiality. On such views, Keller objects that Bernard should save his wife "with thoughts of being true to himself or of retaining his own identity" (2013, 42). But the worry does not apply, at least to my view and possibly to some of his original targets as well. Bernard's love for Patricia is not in the content of any motivating attitude; it is not a thought that moves him. It is rather part of his character; it is why the thought that does move him—"Patricia is in need"—is so powerfully motivating. The courageous person does not act so as to preserve her

partner, child, etc..., as special—of weighing the reasons having to do with them more—is not an attitude about a relationship. It is constitutive of being in a relationship; it is the part of the agent’s psychology that makes it true that she really is a friend, partner or parent. A diachronic requirement to be constant with respect to those principles is a way for a relationship to make a rational difference directly. The very fact that Bernard’s practical reasoning has been that of a partner is what it the case that Patricia’s need matters more. And to emphasize, what is wrong in the case of Bernard’s irrational impartiality isn’t a failure to accord due weight to a relationship, it is a failure to accord due weight to Patricia’s needs, a non-relational fact. What were once on a par—Patricia’s need and the stranger’s—are no longer on a par because Bernard has a history of choosing one over the other. In other words, a requirement to be constant effectively strengthens the weight of the reasons having to do with Patricia. In a way, this view accomplishes the same thing as Keller’s modifier view. But by showing the weights are modified via a diachronic requirement on practical principles we can keep the relationship out of the loving agent’s thoughts.<sup>22</sup> We have found a way for them to play the role relationship theorists wanted while maintaining the moral psychology of individuals theorists.

Nor does the view suffer from the circularity of a synchronic appeal to practical principles. That Bernard may and must count Patricia’s needs more strongly than strangers is not justified by the fact that he does do so, but that he has done so. And lest one worry that

---

identity or to be true to herself and neither does Bernard.

<sup>22</sup> Those who resist the tight connection between reasons and rationality can translate the rest of my view into terms they would accept so long as it is (im)permissible to do what it is (ir)rational to do under conditions of relevant true belief.



this just trades a problem of circularity for one of a regress, consider how such partiality would evolve over time. On first meeting Patricia, Bernard is permitted to treat her as he would any stranger. While it would be wrong for him to save her over  $n$  strangers, he can allocate his attention towards her in just the same way anyone would be permitted towards one stranger among many. That is, before there is a history of being partial to her, he can choose her in situations when anyone could, i.e., situations like our 1 vs. 1 case. Having chosen her, however, he now has a principle that exerts a further rational influence; his history of so choosing makes it permissible to be still more partial to her and, eventually, sets upon him some partial obligations.

The picture that emerges is a conservative one. While a defeasible requirement of constancy does not bind us to our past, it does hold that with time comes a rational pressure to stick to the approaches to decision-making that we've already adopted. If practical principles make up our practical identity—if they are part of what makes us who we are—<sup>23</sup>then this just says that we ought to preserve our identities. As G.A. Cohen would have it, we ought to honor and respect the way we do things exactly because that is who we are (2011).<sup>24</sup> Having gone down a path, it is rational to stick to it; taking some other path through life is no longer as open an option as it once was.

My argument for constancy as a rational requirement has been two-fold: first, it alone of the views considered has squared the circle, achieving the desiderata of both the moral psychology and the ethics

---

<sup>23</sup> cf. Korsgaard 1996.

<sup>24</sup> In the language of Cohen 2011, love often involves “Personal Valuing” (loving an object, be it a person or an old eraser, for its relationship to you), and can involve valuing an object for its “Particular Value.” To be constant is in part to value oneself particularly—to exhibit a bias in favor of one’s extant principles of practical reasoning because that is who one is.

of partiality. Second, inconstancy seems like a good description of what is wrong with a failure of love. But there is a third argument I can offer: adopting this view lets us shed light on deep puzzles about decision-making under parity more broadly. Constancy is not just a virtue with respect to loving practical principles, but practical principles more generally. Having given that argument, I'll then turn to two objections that apply equally to the generalized view and the particular case of love: that seeing constancy as a virtue would allow an agent to illicitly bootstrap her way into making something otherwise impermissible choices rational; and that it is an overly restrictive view that condemns acceptable cases of changing one's mind as irrational.

## 2.6 Constancy as a General Requirement

Suppose Brian is enrolled in a college with an excellent music education program. He can take lessons on any instrument he wants. What's more, to encourage students to try new instruments, the program allows them to switch instruments at any time. There's just one catch: lessons take one week to schedule, so a switch in instruments delays lessons for 7 days. Brian is interested in the trombone and piano. The power and richness of low brass, the fun of the trombone's slide, the instrument's portability, and the opportunity to play with others in ensembles all make the trombone a tempting choice. By contrast, the versatility of the piano, the vast solo repertoire, and the player's ability to make far more complex music on one's own are all reasons to take piano. Suppose all these reasons put the two options on a par, and further that what Brian wants is to achieve an at least moderate competency on an instrument before graduating—splitting

his time in half would be worse than taking four years on just one instrument. At the same time, the parity of reasons resists minor sweetening: if piano lessons had to start a week or two later than trombone lessons, the scales would not be tipped.<sup>25</sup>

Decide Brian must, so on September 1, he chooses trombone, scheduling his first lesson on the 8th. In so doing, Brian will not just have chosen an instrument but also adopted a practical principle. In the face of parity, he weighs pro-trombone reasons more greatly than the pro-piano reasons and begins to act in light of those reasons, scheduling a lesson, reserving an instrument to borrow from the school, etc.... But now imagine that on the 7th of September, the day before lessons are to begin, Brian has a change of heart. Though he had decided on trombone, he gives up on that decision and reopens the question of what instrument to take, sitting down and reasoning anew. Suppose he does not do so on the basis of any new information he learned about the instruments or about his own preferences. He hasn't hit on any reason to think that he made a mistake in his choice on the September 1. Clear-eyed that little has changed and that his first decision was a good one, he nevertheless begins deliberation from scratch and finds that, as before, the reasons for each instrument are on a par with one another. True, his first trombone lesson would be in a day, whereas the first piano lesson would have to be a week later, but that doesn't make trombone the better choice. Brian changes his mind and adopts a new principles of favoring the pro-piano reasons, signing up for his first lesson on the 14th.

Brian has screwed up. While taking piano lessons starting on the 14th is on a par with taking trombone on the 8th, it is worse than

---

<sup>25</sup> See n. 17.

having signed up for piano lessons in the first place and starting on them on the 8th. But while he has made things worse, it is hard to see on a wholly synchronic picture of rationality when he made a mistake. After all, at each of the two times in which he made a decision, the choice between piano and trombone was on a par. So when did Brian make a mistake? It seems like he erred without there being a moment at which he erred.

In light of this, Hedden concludes that Brian did nothing irrational at all (2015a). His decisions constitute a kind of “tragedy,” but not any irrationality.<sup>26</sup> And if rationality were a wholly synchronic matter, Hedden would be right. But if we generalize the earlier constancy requirement, we can say the following: it was Brian’s inconstancy that was the problem in choosing piano on the 7th. Were his previous decision irrelevant, the options would be on a par. But it isn’t. Having adopted a practical principle, he needed some reason to abandon it. Absent a realization that he really did prefer piano or that lessons would be more fun or that he had some reason to think he had erred, abandoning that principle was irrational.

And in identifying the inconstancy as the problem, we make room to say that once he was deliberating anew between between trombone lessons on the 8th and piano lessons on the 14th, the choices were again on a par. It was only the inconstancy—giving up the original principle—that constituted a mistake. Once he’s given up the pro-trombone principle, the damage is done; in that moment, as he begins reasoning anew, choosing trombone tomorrow or piano a week later really are on a par. By locating the mistake in Brian’s manifestation of inconstancy, we can make sense of competing desiderata: that

---

<sup>26</sup> In his (2015c) and (2015b), Hedden defends synchronism about rationality.

on the 7th, piano and trombone are on a par, but choosing trombone and then switching to piano is irrational. It is irrational not because choosing piano or trombone on the 7th is irrational but because giving up the practical principles that makes way for that second choice is irrational. (Another view that can capture these same desiderata is Chang's, which holds that so long as Brian has decided on trombone, his act of will grounds the existence of an additional "will-based" reason to take trombone, thereby breaking the parity between trombone and piano (2013); I'll discuss this view further below in §2.7.1.)

To further the argument, imagine a variation on the case: Brian decides on September 1 to take trombone. But every day thereafter, he gives up on his decision and starts deliberating anew, and each time, picks trombone again. He may not have had to reschedule any lessons, but he still seems irrational. Were we close friends, we might (gently) criticize his lack of stick-to-itiveness and encourage him to be a little more decisive.

People, musical instruments, careers, ways of living—many of the decisions that matter most in life are decisions between options that are on a par. Having decided, those options remain in some sense on a par, viz., so far as synchronic rationality is concerned, they are on a par. But one's past is not rationally irrelevant, and we need some reason to justify abandoning it. We are diachronically extended entities, and our agency can be evaluated not just at a time but across it. Rationality is not just a matter of correctly responding to the world as it is, but to have that response be in keeping with one's past. And though we cannot literally bind ourselves to our previous ways of reasoning, it is a virtue when we keep to them.<sup>27</sup> To fail to do

---

<sup>27</sup> And this is the mistake in Hedden 2015a, which makes too much of the fact that self-binding is, as a matter of empirical fact, often impossible. The issue isn't

so is to alienate yourself from your past, to make oneself out to be a succession of time slices instead of an enduring agent.<sup>28</sup>

## 2.7 Details and Objections

I've claimed that constancy is a virtue. But questions remain: how strong of a reason do we need to justify giving up a practical principle? Would a requirement of constancy condemn as irrational what are in fact normal changes of heart or mind? And does the view fall prey to worries about an illicit bootstrapping, i.e., does my view predict that an agent who makes the wrong choice or adopts a bad principle should stick to it on pain of inconstancy? Let's take each question in turn.

### 2.7.1 Depth and History

How strong a defeasible requirement we have to be constant is a function of two things. The first I've already introduced: the length of time an agent has with a principle. Bernard treating Patricia and the stranger as equally important to him after a few weeks is less objectionable than after a few years. Second is the depth with which a principle is incorporated into an agent's practical identity or sense of self. Brian's favoring of trombone is not at all deep; Bernard's love for Patricia is. What depth amounts to will depend on orthogonal questions in moral psychology. It might be the strength with which one legislates a maxim to oneself, the strength with which you identify with it, the robustness of the disposition to manifest across various

---

whether we can bind ourselves, but whether we should act in accordance with our past principles. We don't need to be bound in order to do that.

<sup>28</sup>cf. Korsgaard 2009.

possibilities, or the importance with which you set an end for yourself. But whatever the details are, the rough idea is clear: some practical principles are more central to us and our outlook on the world.

We can get a theory-neutral sense of depth. Consider careers. The principle to take a certain kind of career-related reason as stronger may be a more or less central part of one's practical identity. Suppose Henry and Carla are graduate students in philosophy; Henry is rather hesitant about the whole thing while Carla is committed. Carla organizes her life around the reasons that motivate and govern philosophers (as opposed to, primary school teachers, doctors, etc...), while Henry might well jump ship after the master's degree. If both are offered a full ride at some law school, Carla's reasons to stay will be far greater than Henry's. Or consider a case of love. Setting questions of time aside, some of our relationships are more central to us than others. The love we've had for an old, but not all that close, friend might be far less deep than a newer, profound connection. Comparing two relationships of similar duration but different depths, there is more rational pressure to be constant the deeper the love. Put another way, the deeper the love, the stronger a reason it would take to make abandoning that love rational.

You might object that these examples are just playing on the fact that we often expect greater costs to be associated with change the longer one has done something or the more invested one is in it. That is, you might concede that length of time makes a difference, but deny that it plays the special rational role that I am claiming it does, instead claiming it only affect us in more quotidian ways. It might be more painful or otherwise harder to find a new career or romantic partner at after a few years than a few weeks, and in that way, history

could appear to be making a difference. A similar objection could be pressed regarding depth. But the problem with this kind of response is by now a familiar one: if these costs were what made the normative difference, then they would have to be reflected in a virtuous agent's motivating reasons. But to continue acting lovingly to your friends in part because finding new friends is hard is a far cry from acting well; and while it would not be inappropriate to be moved to stick with a career or personal project for such reasons, they seem at least to be one thought more than is required. Moreover, such an approach would fail to explain why Brian should not switch instruments in the example in §2.6.

With depth on the table, we might wonder whether it does all the normative work. A longer history often brings with it greater depth, so does history on its own make a difference? Chang's view that our willing creates new "will-based" reasons to do what we have chosen is just such an account; it is just an act of the will that makes a rational difference, not the history of that act. Likewise, Pettit's view of love on which the fact that one loves, but not the history of that love, justifies partiality, recognizes depth but not history. But consider the case of Brian; surely he would be more irrational to switch two or three weeks after his initial choice of instrument, even if the costs associated with extra time (less time on the new instrument) are not enough to break the parity between trombone and piano on their own. The longer its been, it seems the better the reason he would need not to seem irrational in his inconstancy, i.e., our history with a principle makes a rational difference. Or contrast two versions of Bernard. Bernard<sub>old</sub> has loved Patricia for two years where his counterpart, Bernard<sub>young</sub> has loved Patricia for only two weeks. Suppose,



moreover, that the younger relationship proceeded quickly, so that the two Bernards each love Patricia equally as deeply (or that whatever act of will generates reasons is the same for the two Bernards). And lastly, suppose each of them is offered a dream job that would require them to move away, that Patricia cannot move and that as a matter of psychological fact, long-distance will not go well. Intuitively, the job could be good enough (i.e., provide a reason just strong enough) to justify Bernard<sub>young</sub>'s departure but not Bernard<sub>old</sub>'s; where Bernard<sub>old</sub>'s departure violates his obligations to put Patricia first, Bernard<sub>young</sub>'s might not. And, as above, these intuitions about the case cannot be explained by an appeal to the costs associated with a longer relationship or the risks that the younger relationship might not pan out, as those are not reasons either Bernard need act on. A view, like Chang's, that fails to honor this essentially historical nature of these requirements cannot be developed to adequately capture the ethics of love, not to mention its moral psychology (2013).

### 2.7.2 Getting stuck and Bootstrapping

Does a rational requirement to be constant condemn changes of mind or heart that seem rational? To bring out the objection, let's turn to our central case of love. One might worry that a view on which constancy is required would mean we could never end a loving relationship. Suppose that Bernard and Patricia's relationship takes a turn for the worse. Does constancy require they stay? No, and for two reasons. Firstly, there is an important distinction between being "in a relationship" in the ordinary sense of the phrase and the sense in which we ought to be constant in our love. Pace Jeske, it is obvious that we can maintain our love for someone without wanting

to spend time with or interact with that person (2008). Romantic and friendly relationships often end with the parties needing to end their “relationship” and cut off communication and interaction while nevertheless retaining some love for the other. It might be that in some ends of “relationships” we ought to continue loving our former friends or partners, even if such love must not manifest itself in the kinds of action it used to. Secondly, constancy does not yield that you may never stop loving, only that you have defeasible (though possibly strong) reasons to be constant in your love. If that love is too painful or if, as is sometimes the case, we have strong reason to close our hearts to those we once loved, our reasons of constancy are overridden.

A related worry: would a requirement of constancy still render our lives too boring on pain of irrationality? Would we be stuck living life just one way? Granting that we can have a reason to, say, change careers or hobbies or our social lives, does my view require too much of a reason? Again, I doubt it. For one, we can often have very good reasons to be spontaneous or change things up; most notably, the facts that our preferences have changed (“you know, I’d like to try something new”) and that change is good are often sufficient reasons to give up a practical principle and start reasoning in a new way.

Moreover, some of our practical principles are adopted with something of an explicit time limit built in. Brian might decide to study trombone while he is in college; Henry might study philosophy just until he gets his master’s degree and then reevaluate; Carla might be a professional philosopher until she retires. In all of these cases, changing one’s mind is not a manifestation of inconstancy; it does not involve abandoning a principle so much as following a principle’s

built in expiration date. Love, of course, is different. One cannot love just for the season or until I'm 67. One might decide to be in a "relationship" (in the sense of having regular contact with someone) for just a given amount of time, but one cannot build in to one's friendship or partnership a time-limit. That just wouldn't be love. To put the point contrapositively: if I am in a relationship that does have a time-limit built into it (e.g., a summer romance), and I then fall in love, I've screwed things up! The time-limited relationship doesn't admit of love—the relationship either has to change (dramatically) in character or perish.<sup>29</sup> It is not just a feature of western marriage contracts that couples commit to each other 'til death do them part. That reflects the nature of love itself: while it may come to an end, it knows no end, and so to commit to love is to form a commitment that is without end.

What about bootstrapping? If rationality requires that we be constant, does it follow that we should stick to our guns, even when we've erred? Do we have reasons to be constant in vicious principles of practical reasoning—e.g., to be cruel or to be partial only to those with white skin if we've started to be? No. In the first place, constancy requires only that we have some good reason to give up a practical principle. But we of course have very good reasons to give up any vicious or otherwise irrational principle—that just follows from the principle's being irrational. But more needs to be said; after all, we don't want to say that while the racist of 30 years and one of 5 years each have reason to stop being racist, the former at least has a greater countervailing reason to remain a racist. Better to say he has no reason at all. And that, too, is consistent with holding constancy

---

<sup>29</sup> Thanks to Richard Holton for the example of a summer romance.

as a rational requirement. The arguments that got all this going, both about partiality and about practical principles in general, begin with reasoning in the face of parity—between people, between musical instruments, etc.... We needed to explain what was wrong with switching from one option to another where the switched-to option was (in some sense) on a par with the abandoned option. But racism was never on a par with anti-racism, nor cruelty with kindness. The requirement defended has always been one to be constant in principles rationally adopted in the face parity, not to be constant in one's practical principles regardless of their normative status.

## 2.8 Conclusion

By default, our practical situation is filled with parity. Bernard and Patricia might allocate their limited time, attention and resources to anyone, and the reasons to do so are (all else being equal) on a par. But as they choose to favor one another over others, all else is no longer equal. Their love for one another involves a practical principle to favor the other, and the deeper that principles, and the greater its history, the greater the permissions and obligations of partiality they face. By appealing to history we can capture the central role that relationships play without locating them in the content of any agent's attitudes; it is not the fact of the relationship that makes a difference, but the history of so relating that justifies and requires partiality.

I have said much about the ethics of partiality without seriously tackling the question of how partial we may or ought to be. I've tried to illuminate the general structure within which Bernard can save his wife over the lives of a handful of strangers, but what is he to do

if on the other side of the pier are twenty or thirty or one hundred strangers? Indeed, as we are all now aware, the many small decisions we make in our day to day lives—to take a friend to dinner, to go to the movies, to take even a modest vacation—reflect an implicit weighing of benefits to ourselves and loved ones vis-a-vis strangers that is profound.

But while a full account of the extent to which partiality is permissible and required is the topic for its own paper,<sup>30</sup> this picture of partiality at least points the way towards an answer. We took partiality seriously because it was constitutive of love; part of what it is to love is to see one's beloveds (and their wellbeing and preferences) as having a special status in one's practical deliberation. And love is not a practical mistake; a life without love for particular individuals is one that is profoundly lacking, whatever other goods it might instantiate. And so some kind of partiality is part of what it is to live well. How much partiality? At least as much as is constitutive of virtuous, loving practical principles. But to say that our loved ones count for more than strangers is not to give us *carte blanche* with respect to them. While seeing one's son as just one person among many whose interests need to be carefully weighed up against those of strangers is inconsistent with loving him, donating significant portions of one's income to charities and being more modest in one's parental spending habits is not. To be especially crude about it, it seems like while spending 70-80% of one's income on strangers might interfere with loving parenting in a profound way, spending 10% or 30% or even 50% might not be. The arguments of this paper provide a rough, upper limit on the impartialist demands strangers make on us, but it

---

<sup>30</sup> see my MS.

is entirely consistent with a thoroughly partialist ethic that one still owes a great deal to strangers. Put another way, it is consistent with love for particular individuals that one also cultivate, at least to some extent, a love for humanity in general.

## Chapter 3

### Honesty and Discretion

Why should we be honest? Why are lying and deception wrong? And when, if ever, is deception permitted? Disagreement over these questions is the stuff of introductory ethics, where lying and deception serve as focal points in the debate between consequentialists and their opponents.

But for all their important differences, Western philosophers have shared two presuppositions about the topic. The first is about the explanandum, which is taken to be the wrong of lying (or sometimes more broadly of deception) in isolation from other forms of communicative wrongdoing. I'll argue that this "isolationism" leads us to ask the wrong question. Taking lying or deception in isolation from the inseparable vice of indiscretion is a mistake. Neither the scope nor the grounds of our obligation to speak the truth—honesty—can be understood in isolation from its twin, our obligation to respect that which should not be said—discretion. We should not ask "when and why are lying or deception wrong," but more broadly "about what should we be truthful and why?"—a question that is at once about lying, deception and disclosure. Even further, we must at the very

same time ask “about what should we remain silent and why?” These questions, like honesty and discretion, are complements of one another, and answering them together constitutes an “integrationist” conception of the explanandum.

Once we’ve expanded our inquiry to encompass not just the scope and ground of the obligation not to lie or deceive but of honesty and discretion together, we’ll see that a second presupposition must likewise be given up. Western philosophers have approached their shared target with a shared methodology: trying to understand the wrong of lying by considering communicators in the abstract—as speakers and listeners as such, not as friends, teachers, students, or strangers. Many appeal to a global practice of communication to which all are committed qua speakers and argue that lying undermines that practice; some claim that we have a right qua autonomous beings not to be lied to. While special relationships may make a difference in the form of some exception to a general rule, the core requirements of honesty are not seen as depending on local particularities. I’ll argue that this “globalism” should give way to “localism”; instead of starting with global facts about us as agents or communicators, we must begin our inquiry with the local communicative practices of relationships—between parent and child, public official and citizen, teacher and student, etc...—and adopt a relationships-first explanation of the norms honesty and discretion.

I defend an integrated, localist account of communicative virtue—of honesty and discretion. It is a relationships-based view whose main claim is that the norms of honesty and discretion are those constitutive of the relationships we inhabit; they govern what information is “in” or “out of bounds,” i.e., what must and may be shared, what



must and may be kept private, and under what conditions. These norms follow from the nature of the relationships—as intimate, distant, egalitarian, hierarchical—and are shaped by the participants of a relationship in their daily behavior. Thus, the contours of honesty and discretion vary from one relationship to the next. Of course, not all relationships are good ones; some have norms that should be changed or ignored altogether. To address these issues, I’ll defend two supporting claims: that it is insofar as we have reason to be in our relationships that we have reason to follow their norms; and that we often have reason to speak truthfully or to conceal not because the norms of our relationships demand it, but in order to shape those relationships into something new and better.

I’ll begin with two dominant lines of thought on lying in the Western tradition, arguing that their core insights are obscured by isolationism and globalism. I’ll then develop an integrated conception of communicative virtue by taking as a point of departure the local, relational practices that structure our daily lives. Throughout, I’ll largely drop talk of lying for more general talk of deception.<sup>1</sup> I’ll close with a discussion of lying in particular and what is and isn’t morally significant about lying per se.

---

<sup>1</sup>Strictly speaking, deception may not be a genus of which lying is a species. While this is a matter of some debate, so-called “bald-faced lies” in which no deception is intended are of a linguistic kind with standard lies. But we can ignore that complication for now. See n.31.

### 3.1 The Autonomy and Practice Arguments: Shifting the Approach

Notwithstanding important disagreements and variations on themes, Western philosophers have gravitated towards two main kinds of argument about lying. The first, explored by deontologists of various stripes, is an appeal to autonomy: lying and deception violate autonomy and so are impermissible. The second, endorsed either implicitly or explicitly by authors of almost every ethical tradition, is an appeal to the trust necessary for the practice of communication: lying and deception undermine that trust and so are impermissible.

Schematically, the autonomy argument begins with the claim that we are capable of autonomous choice, a capacity which demands respect and can be undermined in various ways. Among those ways is by ignorance—not knowing enough about the context or effect of one’s action can render it less than fully autonomous. To deceive another is to thereby impinge on her ability to choose and to violate her autonomy; therefore, deception is wrong. This basic idea finds very different modes of expression, but something of it can be seen in Aristotle, Constant, Kant, contemporary Kantians, and others.<sup>2</sup>

The practice argument is far more widespread and turns on the significance of trust for communication. In its general form, it begins with three premises: first, that we are justifiably committed to the practice of communication; second, that the practice of communication requires trust; and third, that trust is undermined by de-

---

<sup>2</sup>See Kant (1993, 1998a, 2001, 2017); Constant as cited in Kant (1993); Grotius (2012, §3.1.11.1); MacIntyre (1994); Korsgaard (1986); Faulkner (2007); Rees (2014); Shiffrin (2014); Pallikkathayil (2017); Bok (1999). For Aristotle, see Zembaty (1993).

ception. From these premises, it is argued that deception is wrong. The premises are understood and defended in importantly different ways. Justifiable commitment could be a matter of implicit promising, tacit agreement, hypothetical agreement, common sense understanding, human need or divine purpose. The need for trust is often taken for granted, but can be defended on metasemantic grounds or based on empirical facts about human psychology and language.<sup>3</sup> Undermining might be causal, as for the utilitarians, or teleological, as for Catholics and Kant, or practical under conditions of universalization, as with Kant and most contemporary Kantians. And how those premises yield the conclusion that one ought not to deceive or lie to others will depend on the background moral machinery at work. But these important differences in details notwithstanding, this schema enjoys widespread support.<sup>4</sup>

### 3.1.1 The Explanandum: Motivating Integrationism

As developed, both of these arguments are isolationist, trying to explain the wrong of deception in isolation from other forms of communicative malpractice. The closest related wrongdoing that we can find in the literature is failure of disclosure, sometimes awkwardly called a “lie of omission,” reflecting its status as both like and unlike the supposedly core case of lying.<sup>5</sup> Not as present are discussions of cases of

---

<sup>3</sup>See Lewis (1983); Stenius (1967).

<sup>4</sup>In addition to all of those cited in n.2, see Augustine (1887a,b); Thomas Aquinas (2016, II.II.Q110); Hobbes (1994, §§1.4.3,1.4.12,1.14ff.), Hutcheson (1755, §2.2.10); Hume (2007, §§3.2.1-3.2.2); Bentham (1907, §§16.2.20-24); Mill (2001, pp. 22-23); Sidgwick (1981, §3.7.2, pp.315-16); Ross (1930; 1939, pp. 22-28 & 112-20, respectively); Perry (1909); Williams (2002); Adler (1997); Strudler (2010); Saul (2012a,b); Webber (2013).

<sup>5</sup>See Mahon (2015, §1.1), as well as Mahon (2003); Griffiths (2004); Dynel (2011).

prying or over-sharing.<sup>6</sup> I'll here argue that deception, failures of disclosure, prying and over-sharing should be treated together. The first two kinds of cases constitute modes of dishonesty, the latter two of indiscretion, and we cannot understand one apart from the other. There really is one central virtue (or duty or obligation)<sup>7</sup> when it comes to communication; honesty (which includes recognition of one's obligations to disclose) and discretion are unified in the traditional sense that one cannot understand one virtue (or its correlative vice) apart from the other.

By discretion I mean the virtue of respecting that which should be left unsaid, refraining from both sharing and eliciting information as appropriate. Where dishonesty involves an illicit withholding of the truth, indiscretion is a failure to recognize when the truth ought to be withheld, either by oneself or another. As I see it, honesty is the virtue of speaking truthfully when the truth is called for—not lying, not deceiving, and disclosing appropriately; discretion is the virtue of silence—of keeping silent oneself and respecting the silence of others.

Why go in for the integrationist approach? Most importantly, because it is more explanatorily unified; inasmuch as these two parts of communicative virtue can be illuminatingly explained together, they should be. So the main argument for integrationism just is the articulation and defense of the view that follows in §§3.2–3.3.

But consideration of some cases can help motivate the idea that isolating deception from other communicative wrongdoing is a mistake. To be clear, these cases are not meant as a knock-down argument against isolationism so much as they are meant to provide some

---

<sup>6</sup>Notable exceptions include Nagel (1998) and [redacted].

<sup>7</sup>I will write in terms of virtue, but nothing essential to the view turns on a virtue-theoretic framework as opposed to one which deploys some other normative category.

reason to pursue an integrated account. I'll first argue that deception and failure of disclosure should be treated together, then motivate the inclusion of indiscretion in the desideratum.

**Cancer–Coworkers** Claudia is suffering from cancer but doesn't want her colleagues at work to know about it. Her nosy office-mate, Nick, recognizes some of the visible symptoms. Nick and Claudia aren't close, but he nevertheless asks, "Do you have cancer?" Feeling that it is none of Nick's business and that she isn't ready for the office to know, Claudia wants to get out of this situation. But she recognizes that Nick is perceptive. If she tries to dodge the question or stay silent, he'll rightly conclude she's ill. So she lies with a simple "No."

Claudia's lie is permissible. Note this result does not depend on some intentional wrongdoing by Nick. We can vary the case and imagine Nick as (blamelessly) socially inept and or ignorant of the way in which his question is indiscreet. The result still stands. Contrast:

**Cancer–Partners** Claudia is suffering from cancer but doesn't want anyone to know about it. Her romantic partner, Paige, recognizes some of the visible symptoms. Claudia and Paige are very close, and Paige asks, "Claudia, do you have cancer?" Wanting to keep her illness to herself and feeling that she isn't ready for Paige to know, Claudia wants to get out of this situation. But she recognizes that Paige is perceptive. If she tries to dodge the question or stay silent, Paige will rightly conclude she's ill. So she lies with a simple "No."

Given natural assumptions about their romantic relationship, Claudia does wrong. And lastly:

Cancer–Staying Silent Claudia has no visible signs of cancer and has managed to keep her illness a secret from everyone but her doctor. No one has asked about her health and she has volunteered nothing, neither to her coworkers nor to her partner, Paige.

She wrongs her partner but not her coworker.

That deception and failure of disclosure need to be treated together is illustrated by cases like these. In Cancer–Coworkers and Cancer–Partners, Claudia conceals the same fact from her coworker and partner, but the lie is permissible in the one case where it isn't in the other; so, too, is her silence with coworkers permissible and with her partner dishonest. Because of her relationship with Paige and some kind of attendant norm requiring disclosure, the same kind of communicative act has different moral status. An outright lie can be permissible where non-disclosure about the same fact is a grave wrong. The wrongs of deception and failure of disclosure overlap. What honesty really requires, what deception is permissible—these depend on what norms of disclosure are at play. So the ethics of deception is sensitive to that of disclosure; these two, at least, must be treated together.

Moreover, the wrong felt by Paige in Cancer–Staying Silent is of a kind with the wrong she would feel when lied to in Cancer–Partners. Semantic mistake though it may be, we can easily imagine her feeling lied to in both cases and not being too far off from the truth; no doubt she could feel deceived. That the felt wrong is continuous suggests there is something in common with the wrongs of deception and non-disclosure.

These cases not only show the need to treat deception and non-disclosure together; they also point to the centrality of indiscretion.

Honesty and discretion are often recognized as related, perhaps most famously by Bok, whose book *Secrets* serves as something of a sequel to her *Lying* (1989; 1999). But even she treats the two in separate books, claiming that while a need for secrecy can justify deception, the latter can be understood in isolation.

But honesty is the obverse of discretion; they are inseparable. For one, permissible dishonesty is often a response to (and correction of) indiscretion. Nick's prying in *Cancer-Coworkers* is a violation of Claudia's autonomy just as Claudia's (silent or direct) deception is a violation of Paige's. Nick had no right to the information he sought—it was none of his business—whereas that very fact was entirely Paige's business. Dishonesty and indiscretion are each failures to respect a line between that which is “in bounds” and that which is “out of bounds” to participants in a conversation. Put from the perspective of honesty, it is a line between that which should or should not be shared; put from the perspective of disclosure, a line between that which should or should not be kept concealed. To deceive or fail to disclose—to be dishonest—is to treat a fact as out of bounds when it isn't; to treat a fact that is out of bounds as though in is to be indiscreet.

Nick's was one kind of indiscretion, but we should recognize another:

Over-sharing Ingrid is another coworker of Claudia's, and like Nick, Ingrid is not especially close to Claudia. While they are cordial at work, they have kept to the office's professional norms and never had an intimate personal conversation. Claudia, however, is desperate for some kind of help in dealing with her cancer and unloads everything that she is feeling onto her coworker, from

her fears of death, to her inability to work, to the fact that she is keeping her cancer from her partner. Ingrid is deeply uncomfortable; she wonders how she should respond and whether she should tell their boss or Claudia's partner about what's going on. She wishes she had been left alone; having been told, she must now decide either to keep Claudia's secrets or to share them and that is not a choice she wants to face. Moreover, she knows she has trouble keeping secrets and so now knows she must avoid Nosy Nick around the office.

It could be objected that the wrong here really is quite different from that of deception. What's wrong in this case is just that Claudia has made Ingrid uncomfortable; what's wrong with deception is that the deceiver violates the deceived's autonomy. But that would ignore the reason for which Ingrid is uncomfortable, the underlying wrong to which that discomfort is a warranted response: the fact that Claudia's sharing that which should be left unsaid imposes various (unwanted) constraints on Ingrid's autonomy. Claudia forced a difficult choice onto her, viz. to divulge or to help keep this secret, and constrains the way she can interact with Nick. Again, this is not meant to be a knock-down argument. But to treat over-sharing alongside deception is to recognize that information can be both empowering and limiting, that it can enable, force and restrict choice, and that others can impact our freedom of choice both by withholding and by sharing information. Over-sharing, just as much as deceit, is a way of running afoul of the line. And here, just as in cases where one is tempted to deceive (or pry), the questions faced by an agent are much the same: "what should I say? What may or must I reveal? What may or must I keep silent about?"



It's not that treating the various communicative wrongs—deception, non-disclosure, prying, over-sharing—is impossible. It's rather that the wrongs involved in all of these are related, running afoul of a line between what is in and out of bounds and thereby impinging on the autonomy of others. And we can offer a theory that treats all of these cases together. What we need is a principled and illuminating account of this “line,” the crossing of which in various ways constitutes communicative vice, and respect for which requires at once truthfulness about what falls on one side and silence for that on the other. So where is the line? And what fixes it?

### 3.1.2 Method: Motivating Localism

With this integrated explanandum in mind, return to the autonomy argument I sketched above. It offers the beginning of an answer: in some way or another, the line is a function of our needs as autonomous agents to lead lives without unwarranted interference from others.

But we can see that the autonomy argument faces a challenge. The challenge is made worse given an integrationist explanandum, but it is present even if the autonomy argument is deployed only to explain the wrong of deception: it is not just the listener (the putative victim of deception) whose autonomy is at stake in an given interaction; the speaker's is as well. And much of the deception worth thinking critically about is just that in which the speaker is trying in some way to protect her own autonomy or that of another who has trusted her with proprietary information. In terms of discretion, we see the same dynamic—multiple agents have their autonomy at stake. Nor should this come as a surprise. People's free choices often collide with one another; the Kantian project of theorizing a politics predicated

on autonomy is not aimed at maximizing autonomy simpliciter, but rather developing a framework within which each is her own master consistent with the equal freedom of others. To think that autonomy generates a right not to be deceived is to miss this crucial point: that considerations of autonomy must always be seen in the context of individuals relating to one another.<sup>8</sup> The real question posed by the autonomy argument is not “What communication does respect for an individual’s autonomy require?” It is rather “What is the appropriate balance of communicative rights and powers between these individuals such that each enjoys control over her life in a world where others enjoy the same?”

One way of answering these questions is embodied in the second line of thought in the Western tradition: the practice argument. Of course, those who have written about deception using this kind of argument have not been trying to answer our question about a line between what is in and out of bounds. They have focused only on one direction in which that line can be breached. But we can imagine an answer in keeping with their approach. For Western philosophers, the wrongness of deception was explained by its (somehow) undermining a global practice of communication. That is, the practice in which we engage qua speaker and listener, abstracted from any local practices of communication qua coworker, friend, stranger, opponent, etc....

This global approach is an initially promising strategy. What is wrong with deception towards a friend is, after all, largely the same thing as what is wrong deception with towards a coworker or stranger (modulo some perhaps extra considerations about special relationships). But this is true only when it is wrong; as the cases of Clau-

---

<sup>8</sup>This idea of the demands of autonomy being worked out in relationships with others is given voice in Pallikkathayil (2017) and especially in Ripstein (2009).

dia's lies and silences show, that depends on the local context of communication. Our demarcating line is sensitive to the relationship of speaker and listener. Moreover, if unlike previous theorists of deception, we seek not just to explain the wrongness of lying but the ethics of communication more broadly, the prospects of starting with the global practice of communication or with us qua abstract speakers look even less promising. After all, obligations of disclosure clearly depend on the kind of local considerations that attend to particular communicative relationships (of teacher, of parent, of coworker, of friend, etc...). So, too, do the requirements of discretion obviously depend on their face on local contexts of communication. Thus Claudia's (over-)sharing a fact with her coworker constitutes indiscretion where that same fact must be shared with her romantic partner.

We need, then, a view of communicative ethics that does not abstract away from local contexts or treat these cases as exceptions to be considered after the fact. With insight I'll draw on below, Palikkathayil writes that "certain kinds of relationships, like the relationship between spouses, seem to affect the reasons we have to [speak the truth]." But she goes on to say that "[e]xamining exactly how and why this is would take us too far afield" (2017). The object of our inquiry, however, is afield. The line, after all, varies locally, so we must understand honesty and discretion through the lens of local relationships. That does not mean we give up on a unifying, explanatory thesis that yields universal truths about honesty and discretion—quite the contrary. To be clear, this is not a defense of relativism or a claim that there is not some underlying set of normative truths governing the ethics of communication in all contexts; it is a claim about method and emphasis: we must start with the local, the particular

relationships that set the context of communication, and build outward. Instead of treating them as exceptions or qualifications to a general rule, these local contexts must be given pride of place in our theorizing.

Moreover, adopting a local approach addresses the biggest problem with virtually all instances of the practice argument: the third premise, that deception undermines trust, is false, at least if it is understood globally. Trust is fine-grained, it admits of nuance and context-sensitivity. What it takes to be trustworthy must be understood locally: trust is undermined not by deception as such but by crossing the line.

To take a representative, contemporary statement of the thought:

[T]here are no alternative, precise, and authoritative avenues into the contents of each other's minds; there is only testimony. To [deceive and thereby] use this avenue of knowledge for a contrary purpose is to render it unreliable and to taint it.... In this case, versatility is a vice. By doing so, one has eliminated a fail-proof, trustworthy mode of access to one another. (Shiffrin 2014, p. 23, emphasis added)

Shiffrin's statement may seem extreme, and she introduces nuanced qualifications to the view that accommodate common views about permissible cases of untruthful speech. Still, her main claim is clear, and it is one that we can find endorsed over and over throughout Western writing on lying.<sup>9</sup> But the claim that versatility is a vice,

---

<sup>9</sup>See Augustine (1887b, §11); Grotius (2012, §3.1.11.1); Hutcheson (1755, §2.2.10); Kant (2001, Ak. 27:444), Perry (1909); Korsgaard (2009). Most rigorists about lying (e.g., Aquinas 2016) seem to at least implicitly endorse this.

that open communication is liable to tainting by any kind of deception, fails to respect the way in which our capacity for trust is fine-grained. Trust is not on-off; trustworthiness must be understood in its local context.

To illustrate the point, consider the carefully crafted system of rules governing truth-telling and concealment in court. As Sidgwick observed in his criticism of Kant, “In the word-contests of the law courts, the lawyer is commonly held to be justified in untruthfulness within strict rules and limits” (Sidgwick, 1981, §3.7.2). Consider:

Defense Attorney Scout, a defense attorney regarded as a woman of great integrity, is defending a client she knows to be guilty. In her opening statement, she tells the jury, “My client is not guilty—he cannot have killed his friend. And over the course of the next few days, I’m going to show you why.” She proceeds to present the evidence of the case in a strictly-speaking truthful manner, but always showing things in a light that warrants the inference that her client is not guilty. In so doing, Scout lied (about her client’s guilt); she then deceived by misleading and misdirecting jurors.

I think it is clear that Scout does no wrong; nor does she wrong (in a directed sense) the “victims” of her many deceptions and single lie. But let me emphasize, the point of this case is not to provide a counter-example to extant views on lying. Shiffrin and others have had sophisticated and insightful things to say about cases like these.<sup>10</sup>

---

<sup>10</sup>Shiffrin takes as central the notion of a “suspended context,” a local zone in which the listener is not warranted in taking testimony as truthful (2014). The law courts would, on her view, be a “justified suspended context,” i.e., one on which the speaker may justifiably speak (somewhat) insincerely. But while her view can accommodate the case, her conception of trust as on or off, “eliminated”

The point of the example is to illustrate a flaw in their conception of trust; they try to see it through a global lens, where it must be understood more locally. Scout's trustworthiness in general is not undermined by her deception. We can imagine the judge in the case following up with Scout, asking her whether she and her client would be able to accommodate a scheduling change, for example, without doubting her trustworthiness on the matter. The judge might even see her as a particularly trustworthy defense attorney, e.g., trusting completely that Scout has obeyed all the rules of discovery, handing over evidence that she might rather have kept to herself. All this while at the same time recognizing that in doing her job, she may lie about her client's guilt and offer misleading arguments that present her client in a false light. The judge's trust in her admits of a nuance: as advocate, Scout is not fully reliable about that which relates to her client's guilt or innocence, but about other things, she is.

In this case, a highly artificial set of rules is what allows the judge to trust Scout in such a fine-grained way; the line is given by the rules of criminal procedure that specify exactly how, in what way, and for what purposes Scout can lie, mislead, and deceive and when she must disclose, answer forthrightly, etc....<sup>11</sup> We have set up a practice that

---

by deception, is too coarse-grained. On some readings of Kant a statement does not count as a lie (and so one has no duty to refrain from making it) if intending the listener to believe the untruth does not thereby violate her rights. On this reading, Kant, too, has a notion like Shiffrin's suspended context, and possibly an even more fine-grained one. Whether it will be subject to the same criticism depends on what rights one has to various kinds of information; if those rights are highly dependent on the relationship between speaker and listener, then Kant's view might count as suitably localist. However, as Mahon makes clear, this is just one of a number of senses of lying in which Kant is interested, and his other senses of lying and the corresponding duties to refrain from lying, are insensitive to the particular rights at stake (2009).

<sup>11</sup>An article in the *Stanford Law Review* claimed, "A lawyer is required to be disingenuous. He is required to make statements as well as arguments which he does not believe in. But the further his statements descend towards the particular,

not only allows for but enshrines in publicized doctrine lying and deception in a way that doesn't undermine trust in general. That is not to say that defense attorneys are all fully trustworthy; of course, many are not. But that isn't just because they lie and mislead. Those who are untrustworthy in general are so because they push or flout the rules or mislead in ways that fall beyond the prescribed limits. They are untrustworthy not because of their deception per se but because of the way in which they cross the line.

This is all to say that effective communication does not require trustworthiness about all things in all ways between all people. We must understand trust locally. It admits of a fair bit of deception provided it occurs within some kind of limits, i.e., with respect for the line. In Scout's case, the limits are an explicit and artificial set of rules, making the permissibility of deception especially easy to understand. But the phenomenon is pervasive and less commonly a product of thoughtful artifice than the organic evolution of communicative practices. Recall Cancer-Coworkers; Claudia does not undermine her general trustworthiness by lying about her health. If the lie comes to light, Nick could (and should) recognize that while Claudia is willing to lie about a matter as private as cancer, he can still trust her in matters of business. Many sports often have an implicit set of rules that govern when athletes can be trusted and when they are deceiving in keeping with the norms of the sport.<sup>12</sup> Some games are predicated

---

the more truthful he may be, indeed must be, because no one appreciates the significance of the particular better than a lawyer.... [W]hen he is talking for his client, a lawyer is absolved from veracity down to a certain point of particularity" (Curtis, 1951).

<sup>12</sup>Professional cyclists often downplay or outright lie about illness during a race if they are in contention and do not want opponents to know about it. Professional soccer players "dive" and "flop," and often defend the practice as an integral part of the sport. Alexi Lalas, a professional turned commentator described diving "...as a skill. There are good dives and bad dives, and the act of embellishing an

on deception, e.g., from poker to the board game Diplomacy, in which the object is to make alliances with other players and betray them. Friends can engage in such localized deception within the game (“she’s probably lying about what she’ll do next turn...”) while still trusting one another in other matters (“she said she wants to reschedule next week because of a doctor’s appointment...”).

To proceed, we must keep two things in mind: first, that our conception of the explanatory target must be integrated: the ethics of communication as a whole. And second that our approach must be suitably local, developing an account of the scope and grounds of communicative virtue (or of our rights and obligations as communicators) that in the first place respects local communicative practices—between friends, coworkers, partners—and their particular needs and norms.

## 3.2 A Relationships View

Instead of thinking of some universal of communication, we should instead recognize that we have many different practices of communication that have different ends and so different standards of truth-telling. Scout has reason to lie about her client’s guilt because she has reason to act as a defense attorney, and what it is to relate to a jury qua defense attorney requires lying about her client’s guilt.

What follows is a revised practice argument. Instead of thinking of the practice of communication as such, we’ll consider what kinds of communicative practices there are, what various ends (like coop-

---

action, I don’t see as an insult to the game or the person.” He went on to complain that attempts to crack down on the practice are “robbing the game of some of its most interesting tactics” (Irfa, 2018).



eration, justice, or emotional intimacy) they might serve, and how different communicative norms serve those ends. We can here take inspiration from various Eastern writings on lying and deception, where in contrast with the Western tradition, Confucius's Analects and the Baghavat Gita offer a local, integrationist approach.<sup>13</sup> They take as their starting point the concrete, social relationship between speaker and listener. And without adopting the strict social hierarchies of these views, I will argue that we should adopt their general orientation towards communicative ethics: to situate it within the ethics of relationships, where where what one should share and conceal is largely a matter of to whom one is talking and the relationship between speaker and listener.

I will argue that relationships have constitutive standards, among them constitutive communicative standards. Then we'll turn to the question of when those standards get a normative grip on us. I'll argue it is when we are (actually) in a relationship and have reason to be in that relationship. Roughly, relationships provide the rules, and those rules provide reasons for us just when we have reason to be in our relationships. Recall the question posed by the autonomy argument (§??): how are we to strike a balance between the conflicting rights and powers of individuals trying to live autonomously together? A relationship constitutes a set of norms that governs just that—it is a local, candidate answer to the question of autonomy. Some of these proposed answers are better than others, and it is the good ones that get a normative grip on us. A picture with relationships at its center can maintain what is right about both the autonomy and

---

<sup>13</sup>Confucius (1997, see especially §13.18, §13.30), *The Mahabarata* (2016, Drona Chapter §§191-93, Karna Chapter §69). See also Bonhoeffer (1965), which offers a relationship-based conception of what it is to tell the truth, and MacIntyre (1994), which recognizes the centrality of relationships to norms of telling the truth.

practice arguments: deception (and other forms of communicative malpractice) are wrong when and because they cross the lines of a given communicative practice—a line which is an attempt to negotiate the question of autonomy between individuals.

But relationships are dynamic, not static practices. Their ends and norms are set by us and can be changed by us. Seeing that relationships are fluid, with communicative norms that often change just because someone does or does not tell the truth, will equip us to paint a layered and powerful view of the ethics of communication. We'll see that we have two distinct kinds of reason to be truthful:

1. Because the norms of a relationship that we are actually and should be in require it.
2. Because doing so can better the relationships we have with others.

To get to all of that, we first need a better sense of what information and communication do. We'll then turn to the roles it serves in relationships and finally put it altogether with an account of when one should follow a relationship's constitutive standards.

### 3.2.1 Information and Communicative Practices

Defenders of the old practice argument really see communication as having one crucial function. Above we saw Shiffrin identify that as the sharing of one's thoughts, an echo of Augustine:

Speech was given to man, not that men might therewith deceive one another, but that one man might make known his thoughts to another. (Augustine, 1887a, §22)

And virtually all on the topic agree: speech is for sharing one's thoughts in the service of getting along with one another.

But openly sharing one's thoughts is hardly a straightforward recipe for sociality. Quite the contrary, getting along with one another frequently requires that we conceal our thoughts—with silence and even deception (Nagel, 1998). The thought of someone blurting out all that he thought or answering truthfully any question asked in social settings (a party, a department meeting) is enough to make one squirm.

If communication isn't solely aimed at sharing one's thoughts, what other functions does it have and why? In this section we'll consider in general the variety of communicative functions and ends; in the section that follows, we'll use that to build our relationships view of communicative virtue.

Sometimes, a given communicative practice will have a very specific purpose. The courts aim at justice and so have rules of communication which serve that aim. Professional journals aim at the advancement of the field and so set up rules to serve that (e.g., state all your sources but conceal your identity during review). Various games and competitions require the concealment of some things but the sharing of others (e.g., a poker player's funds are displayed openly at the table but her hand is hidden). So the first thing to note is that particular human activities regularly do fashion norms of communication to serve their particular ends. Indeed, it is hard to think of an end-driven, group activity that does not have some kind of attendant communicative structure. It would be impossible to survey all the aims to which we put communication and how we create normative structures to accomplish those aims, so for now, I'll say only some-

thing very general about the way that informational structures play out in our personal relationships.

The having and sharing of information shapes dynamics of power, vulnerability and intimacy. Communicating isn't just about sharing our thoughts—its function is also to structure these dynamics between individuals. Consider the upshot of another's knowing something private about you—your medical history, your preferences, a secret of which you are ashamed, etc.... That knowledge gives its owner power: knowing your likes and aversions can give someone the power to humiliate, harm, manipulate, persuade or motivate you. Before a job interview, the interviewee might endeavor to discover his interviewer's preferences so as to better position himself. A doctor must know intimate medical details to treat her patients. These are all instances of an agent gaining an ability or power (to position himself, to cure) in virtue of her knowledge. Correlatively, to have something known about you is to make yourself vulnerable. Sometimes that vulnerability can be a weakness—we keep our medical information private lest we fall victim to those who would abuse us; but it is often desired, as with one's doctor. Sharing one's private life with friends and family makes us emotionally vulnerable; we entrust our hearts to the hands of another. That vulnerability fosters trust; it is constitutive of intimacy. We share private information because that is a part of sharing oneself with those we love and trust.

This power and vulnerability come not just from the possible effects of having information but also from the having in and of itself. Imagine one of your closest held secrets being discovered by someone

you didn't intend to know. You might worry that he could use that secret to ill effect; but even if you knew he would keep it to himself, his very having it is enough to warrant discomfort at being so exposed to another. This power and vulnerability can even come from one's merely having an ability to get some bit of information in future (without yet knowing it). Citizens of a state that fails to respect privacy rights aren't just impinged upon when the state does look into their personal lives—the fact that at any time the state can look is enough; whether the doctor has yet looked into my medical charts, I am vulnerable in virtue of her ability to do so; the fact that my partner can ask about my innermost thoughts is part of what constitutes our intimacy even before she knows the thought itself.<sup>14</sup>

Thus, how much openness of information characterizes the informational dynamic between individuals structures all kinds of things about their relationship. How intimate a friendship is depends in large part on how much information is in bounds in that relationship—on what can be asked about, shared, etc...—and what is out of bounds. The political relationship between citizen and state depends very much on how much information the organs of the state have or can obtain about citizens and how much about the functioning of the state can be known by citizens. The dynamics can be shaped in a fine-grained and context sensitive way. A doctor knows the most intimate details of our medical lives but needn't know anything else; we are highly vulnerable in just one domain of information. A thesis advisor

---

<sup>14</sup>While traditionally, American courts have refused to recognize mere disclosure or private information as a legal injury absent some further effect, a growing body of case law is establishing as a genuine harm another's merely knowing that which should have remained private. (Melvin v. Reid (1931); Coverstone v. Davies (1952); Gill v. Hearst Publishing Co. (1953); Samuel v. Curtis Pub. Co. (1954); Timperley v. Chase Collection Service (1969); Porten v. University of San Francisco (1976))

might know the ins and outs of her advisee's work habits so as to better advise him without knowing much at all about his personal life; there could be great intimacy, trust and vulnerability about the student's intellectual life without extending that vulnerability and trust to other areas.

With a richer understanding of why information matters and how it can shape power, vulnerability and intimacy, let's now put it to work to build a view of honesty and discretion.

### 3.2.2 Relationships & Constitutive Ends

Some practices have constitutive standards of evaluation; that is, performance can be evaluated on the basis of standards that follow from the nature of the practice. A teacher can be evaluated as better or worse qua teacher; the standards of evaluation follow from what it is to teach. Part of what it is to be a teacher is to aim at the education of one's students; that end, or telos, is part of what makes a given practice teaching as opposed to, say, oratorical performance. To oversimplify a bit, a teacher can be evaluated by how well she achieves the end of imparting knowledge and skill in her students. To evaluate her qua teacher, we can ask how well her students grasped the material, learned the skills of the field, etc.... We can make similar evaluative judgments of many other roles that are situated in a practice: being a baseball player (or more particularly a pitcher), being a chess player, being a doctor, being a defense attorney—these are all roles that have ends, and occupants of those roles can be evaluated qua occupant by asking how well they achieved those ends.

Of course, the standards of a practice may or may not be normatively relevant to an agent; to say that an agent is evaluable relative

to some standards is not yet to say that the agent ought to follow the standards. Torturers are better qua torturers the more pain they cause; but that does not give the torturer more reason to use harsher methods. The torturer should not be a torturer at all. It is a separate question when constitutive standards of a practice actually generate reasons for an agent; we'll turn to that question in §3.3.

To claim that a relationship is a practice with constitutive ends is to say that participants in a relationship are evaluable qua members of the relationship. It makes sense to say of someone that he acted well as a friend, poorly as an advisor, etc.... (Again, to say that someone is acting well as a friend is not necessarily to say that they are acting well simpliciter, e.g., "he's being a good friend in spending so much time helping Berk, but at this point, he needs to focus on her own work.") In some relationships, the ends of the participants' roles are, in general outline, obvious: a thesis advisor tries to make her advisee's thesis as good as possible, a doctor keeps her patient healthy, a baseball coach improves his players skills. But what could the end of a friend or of friendship be? And obvious or not, where do these ends come from? What explains why a role has that end?

Asked in a certain tone, the question seems to suggest that any answer is metaphysically suspect; Aristotle might have an idea of the end of friendship, but we are supposed to be skeptical of any such natural teleology. For my own part, I am less than fully skeptical;<sup>15</sup> but we need not rely on any kind of natural teleology to get the picture going. For the most part, our relationships are shaped by us; we set their ends, explicitly or implicitly. This friendship may be about

---

<sup>15</sup>Part of what it is to be a mother or father or guardian is to aim at something that seems naturally given (the flourishing of the child in one's care). Indeed, what it is to be a friend might also be at least in part a function of the kind of social beings we are (Foot, 2001).

having fun and keeping things light when together; that may be about deeper, emotional support, with the friend sharing more of themselves with one another. This dissertation advisor-advisee relationship is about deep collaboration and involved support; that about providing occasional feedback but largely letting the advisee make her own way.

A full metaphysics of relationships is its own project in social ontology. All my account needs is for relationships to be activities that can be characterized at least in part by their function. While it might be somewhat odd to put it in these terms, relationships do clearly serve functions in our lives. Some relationships' ends are largely instrumental: the doctor-patient relationship serves the latter's health, the coworker relationship the success of the business, the coauthor relationship the success of the research, etc...; others are partly instrumental: healthy friendships, partnerships and familial relationships aim at the wellbeing of their participants. And, of course, many relationships are ends in themselves—the character of the interaction between the participants (as loving, joyful, intimate, etc...) is not just a means to something else. That a given relationship has the ends it does is part of what makes the relationship what it is (a friendship, not a relationship of mere collaborators; a romantic partnership, not a casual fling). And while a teleological conception of relationships threatens to be implausible if the ends at issue are assumed to be simple, we can recognize that a given relationship may have many, nuanced and highly particularized ends set in some way by the intentions of the participants, the actions that lead to the relationship itself, and the social environment in which that takes place.

Thus Claudia and Paige's intention to be romantic partners, their mutual commitments (explicit and implicit), and joint decisions to



plan their lives in concert make their relationship an intimate partnership with the ends that it has (emotional and physical intimacy, mutual support, fulfillment, the flourishing of the partnership itself, etc...). Their ends require, both constitutively and as a means, a ceding of individual control. More information is in bounds because these are lives lived together, and the line between what is in and out of bounds must serve these intimate, shared ends. Claudia and Nick have a coworker relationship whose ends are to work together effectively in a professional environment, requiring a greater zone of individual control, informational and otherwise, about personal matters.

In light of these constitutive ends, a participant in a relationship can be judged as a participant in a relationship. Following from the constitutive ends of a relationship are constitutive standards, adherence to which serves the ends of the relationship and define the character of the relationship. Among the standards of a relationship are the obvious: advisors should meet regularly with their students, friends should make time for one another, etc.... That the one person should make time for the other follows from what it is for the one to be the other's friend. And crucially, among the standards will be communicative ones. After all, the ends of a relationship can be better or worse served by different communicative norms: Claudia and Nick's coworker relationship is ill-served by deep emotional intimacy, so much is off-limits; Claudia and Paige's partnership, by contrast, requires much more intimacy. Relationships with identical ends might still have very different constitutive standards on account of the difference in their participants' needs, preferences and decisions. Where one advisor-advisee relationship would benefit from personal as well

as intellectual closeness, another might be thus encumbered.

In general, these standards, along with the ends that they serve, are dynamic, shaped and reshaped by their participants in the very acts of relating to one another. Sometimes this shaping is done explicitly, as is often the case in legally, contractually, or institutionally regulated relationships. An advisor and advisee might discuss the terms of their supervision at the start of each year; a couple in marriage counseling might explicitly discuss and agree to norms, communicative and otherwise, that better serve them. But more often than not, these ends and standards are shaped implicitly, not by talking about behavior but by the doing itself. That a relationship is shifting from that of mere colleagues to one of friendship would be more likely accomplished by behavior than explicit agreement: that you asked me to lunch, that I gladly said yes and reciprocated with a follow-up the next week—we are using these social signals to communicate our intentions to become closer. As we talk, my opening up about life outside work and your taking up that opening (listening warmly, asking questions, not redirecting the conversation) shapes our relationship and so its communicative norms to make life outside of work firmly in bounds. Likewise, redirecting the conversation away from a topic, refusing to answer a question, answering vaguely, these are ways of placing limits and establishing what is out of bounds.

Friendships are often reciprocal and egalitarian; so one participant's sharing a kind of information, say about one's past, and the other's taking it up is often enough to put that (and related) topics in bounds for both participants, all else being equal. (What might not be equal? One participant indicating, explicitly or implicitly, that they'd rather not share but are happy to listen.) Other relationships

lack that reciprocity; when a patient shares private information with his doctor and she takes it up, that in no way entitles the patient to ask for similar information about the doctor. When a public official divulges private information to the public, that in no way entitles her to the same information from the public she serves. Whence the difference? As always, the reciprocity or lack thereof serves the ends of the relationship (of intimate friendship, of a given political order).

What of strangers, those with whom we have no special relationship? Even with them we are in a relationship; we are not born as Hobbesian atoms but always already in relationship to others. Society is structured by norms of behavior that facilitate the end of getting along together. Of course, different societies do so differently, and their norms of communication reflect that. Consider the norms of small talk. In the US, where work is a part of many people's identities and certainly of our national conception of ourselves, asking a stranger what she does for a living is commonplace; in France, with its great commitment to egalitarianism, it is downright rude (Barlow and Nadeau, 2016). In India, where the social structure is imbued with far greater intimacy by default, it would be neither surprising nor invasive to ask someone sitting next to you on a plane about how much money they make nor even to express (dis)approval at the figure (Mahajan, 2016). Or consider the place of talking about politics or religion—whether it is acceptable is in part a function of whether those aspects of a life belong in the public sphere. These are all differences in the norms that govern strangers' interactions; they reflect differences in the aims of a social order and how its inhabitants get along.

At this point we've arrived at the core of the view. Our relation-

ships have ends, and we as participants in those relationships can be evaluated in light of those ends. The norms of evaluation include communicative ones which fix what can and may be shared and in what circumstances, thereby shaping dynamics of power, vulnerability and intimacy in service of the ends. Being guided by these norms constitutes (most of) communicative virtue—at once honesty and discretion (the qualification comes in the following section).

If this is the fundamental picture, we are faced with new, important questions. Where standard accounts of lying must concern themselves with exceptions to their general prohibition, the relationship-based account has these “exceptions” baked into the account. But we face new issues: which norms (i.e., which relationships) ought to govern our behavior? What of relationships that are defective and need to be changed or abandoned? And what kinds of relationships should we seek to establish?

### 3.3 Which Relationships Matter

To a large extent, these questions are specific instances of much more general questions that face two-level normative theories—that is, theories in which particular actions (the lower level) are evaluated not directly but in light of some set or rules, norms, practices, dispositions, etc...(the higher level).<sup>16</sup> Any such account needs a principle that specifies when an element at the higher level (the rule, practice, disposition) actually matters normatively—when the action in question should be evaluated in light of that element. We’ve already seen

---

<sup>16</sup>Familiar views in this family include rule utilitarianism, Contractualism (of various kinds), Contractarianism, Foot’s virtue ethics (2001), Kantian approaches (of Kant and followers). See also Thompson (2008, Part III) and Kagan (2000).

that in our case, the following principle fails:

Transmission-Actual When an agent is in a relationship, she has reason to follow the constitutive standards of that relationship.

Torturers shouldn't torture more painfully, they shouldn't be torturers at all! A woman in a patriarchal marriage isn't required to follow its norms just because she is stuck in it. She has every reason to relate to her husband differently. If her husband asks where's she been or what she's thinking, she can ignore the standards of their relationship and dissemble or lie. She may, of course, have reason to pretend to be in the relationship if her safety requires it; if she lacks options, she may even have reason to stay within the legal institution of marriage and continue pretending to relate as a wife. But she should not treat her marriage's norms as genuinely normative for her.

This suggests an immediate fix:

transmission If an agent is actually in a relationship, then insofar as she has reason to be in it, she has reason to follow the constitutive standards of the relationship.<sup>17,18</sup>

---

<sup>17</sup>Is this an instance of a more general transmission principle about all practices, relational and otherwise, e.g., individual habits? I doubt it, but the arguments to settle matters one way or the other are both complicated and unnecessary for our purposes here.

<sup>18</sup>There are alternative candidate principles, but the issues involved in settling on one are largely orthogonal to our topic. I will flag that a purely hypothetical version of the principle, e.g., insofar as an agent should be in a relationship, she has reason to follow the constitutive standards of that relationship, fails for reasons analogous to non-compliance objections to rule utilitarianism (as in Parfit 2011, pp. 308-320). We can have decisive reason to be in relationships we are not yet in, and those ideal relationships may recommend actions that are inappropriate or disastrous in a context in which that relationship is still a mere ideal. For instance, imagine Anthony and Ayesha are coworkers whose relationship constitutively involves keeping great personal distance. By accident, Anthony has come to learn that Ayesha is in the midst of a difficult divorce. As her coworker, it would be inappropriate for Anthony to ask about how she's handling it; he's not even meant to know. However, Anthony and Ayesha get along quite well.

This principle still raises questions about partially defective relationships, i.e., those we have some reason to be in but some of whose norms seem unacceptable. This objection—akin to that of rule worship against other two-level views—is the topic of the following section (§3.3.1). But first, let's ask what reasons we have to be in a relationship. All kinds! We might promise ("I'll supervise your dissertation"; "I do"), or enter a contract (to be a baseball coach or doctor). Circumstances may provide reasons (having had a child, he has reason to be a father towards her). And most often, it is generated by ordinary behavior. We act in a way that signals our interest in being in such and such a relationship, implying our commitment or at the very least engendering expectations that we will so relate.

And it is here that we should return the autonomy argument. Proponents of the autonomy argument thought that autonomy was central to the wrong of deception and in this respect they were right: the norms that really do make most deception wrong are those that follow from the relationships we should be in, and we should be in relationships that promote the autonomy of their participants! Autonomy is among the key lenses through which relationships themselves must be evaluated. So when evaluating the reasons we have to occupy a given relationship, we should ask how well that (local) practice balances the rights, powers and liberties of its participants, how well it negotiates

---

They would make great friends and have every reason to become friends. Were they friends, Anthony should ask how Ayesha's doing. Anthony should become Ayesha's friend. But he does not yet have reason to act as her friend and ask. They lack the requisite trust for his inquiry to be well received; instead of a caring inquiry from a friend, it would still be a rude and presumptuous intrusion. What Anthony does have is reason to become Ayesha's friend; they get along well and she needs friends now especially. But the appropriate, respectful and kind way to treat someone one would befriend is quite different from the way to treat someone who already is a friend. So we need the principles to require that the relationship be actual, and not merely ideal.

the interaction of their lives.

And because relationships are the way in which we try to find, or rather create, the space for other individuals in our lives, relationships are better when they are not all-consuming. The relationships we have most reason to be in and to create with one another are those that build in a recognition of the limits of the relationship, whose norms reserve a space of autonomy (informational and otherwise) for each individual. Healthy friends, spouses, thesis advisors and coworkers do not control the lives of their counterparts. Communicatively, that means that in the healthy version of these relationships, much is kept private—no participant is required to share all and be wholly communicatively vulnerable and exposed. Healthy relationships honor privacy and so permit some concealment. This is not some unfortunate bug in relationships but a feature, one made especially salient because we have considered relationships with an eye towards our integrated explanandum. Being a coworker would be (untenably) exhausting if everything were in bounds; even many friendships would be unbearable if the friendship left no informational control in the hands of the individual.

It does not follow that deception is to be permitted *carte blanche*. But it does mean that some deception, and maybe even lying, must be allowed. Sometimes we pry—be it accidentally or intentionally—and our interlocutors must have the freedom to resist our intrusion into their space. The ends to which we put communication require it. This isn't the end of honesty. The limits and scope of this permission are narrow, to be called upon only as a last resort; and violation of those limits is a very grave wrong. The point is just that some degree of deception is necessarily permitted by the nature of our relationships.

Claudia must be allowed to deceive, as must Scout—they have secrets to keep, and given the circumstances, the only way they can do that is with some scope of deception. Staying silent, a strategy often recommended by rigorists, sometimes reveals too much.<sup>19</sup> As Bacon put it:

He that will be secret must be a dissembler in some degree. For men...will so beset a man with questions...they will gather as much by his silence, as by his speech.... So that no man can be secret, except he give himself a little scope of dissimulation; which is, as it were, but the skirts or train of secrecy....

The best composition and temperature, is to have openness in face and opinion; secrecy in habit; dissimulation in seasonable use; and a power to feign, if there be no remedy. (2009, §6)

Nor should we be afraid to allow the limited deception that Bacon recommends. We've seen that our practice of communication can tolerate some nuance; and provided that this deception is in some principled way constrained (i.e., by the norms of the relationship), it need not undermine the trust necessary for communication in general. To the contrary, it makes creates the possibility for relationships and roles like that of a non-intimate colleague or of a defense attorney that require some concealment.

---

<sup>19</sup>As in Augustine (1887b, §24); Thomas Aquinas (2016, II.II.Q110); Korsgaard (1986, p. 330, n.4).



### 3.3.1 Flawed Relationships

Of course, not all relationships do this perfectly. Most of our actual relationships are flawed answers to the question of autonomy. What are we to make of relationships whose norms are in some way deficient but which seem nevertheless like relationships worth being in? Recalling the worry about rule worship, what are we to make of relationships that seem to satisfy the antecedent of transmission (we have reason to be in them), but falsify the consequent (they have norms that should not be followed)?

I'll here explore a case of a deficient relationship. My aim will be twofold: I will defend transmission and so the view as a whole, against what I see to be the most important kind of counterexample; and in so doing, I will bring to light the second kind of reason we have to be truthful or to conceal: to shape our relationships.

I want to flag in advance what I am not doing in defending my view: appealing to a Rawlsian notion of a practice (as opposed to a summary) rule, viz., a rule that defines a new move within a practice (Rawls, 1995). While relational practices do have some communicative practice rules ("do not commit perjury" is a notion that cannot be defined without appeals to the communicative norms of the law), I do not maintain that in general, what we are doing when we speak to people can only be understood in terms of the rules of that relationship.<sup>20</sup> Instead, I'll argue that what at first seem like counterexamples to transmission are not yet fully specified. When we do fully resolve one of these cases, it will fall into one of three categories, being either one with unarticulated exceptions to what seemed like a

---

<sup>20</sup>Bonhoeffer (1965) seems to endorse the idea that truthfulness itself is defined only in the context of a given relationship, but I don't know how literally he should be read, nor can I see how to make a literal interpretation work.

rule, one in which the agent should change the relationship, or one in which the agent does indeed have (a perhaps overridden) reason to share. Consider:

In Need of Privacy Charles and Julia's marriage is one defined by its unusual openness. Both having suffered from trust issues earlier in their lives, they share everything; it is a relationship in which neither claims any domain of privacy from one another. Charles starts seeing a therapist and with help identifies that too much of his life revolves around Julia and his marriage; the therapist encourages him to delve into these feelings, start a diary, and figure out what he needs. The therapist urges Charles to embark on at least the initial phases of this without Julia; he recommends that without deceiving her, he nevertheless keep much of this private, at least at first, so that he can pursue the issues independently.

Let's stipulate first that whatever problems there are with this relationship, it is not a wholly unhealthy partnership, and second, that the therapist is correct. Charles (rationally) wants to stay with Julia, while at the same time needing to work some things out without Julia's involvement or knowledge. This threatens to be a counter-example in which Charles has reason to be in the relationship but, contra transmission, no reason to obey its norm of total openness.

If that is the right normative description of the case, transmission is false. I'll argue that it isn't. As stated, the case is underspecified. When Charles shares the therapist's recommendation with Julia, what does she say? Is she supportive? Willing to make adjustments to their relationship? Insistent that he nevertheless share? Her reactions matter for two reasons: one, they might provide new reasons for

or against a course of action; but second, they, along with Charles's feelings, might tell us something about the nuances of their relationship itself. We can further specify the case in three ways; and in each, we'll see that transmission holds.

Version 1: Unarticulated Exceptions, no reason to share

This might be a case in which the norms of the relationship have unarticulated exceptions, where Charles does have reason to be in the relationship, but the relationship does not issue in a reason to share. Most relationships are not maximally strict in their constitutive standards; they admit of unarticulated exceptions from the get go, even (or especially) when the norms are somewhat strict. Suppose you have a strong commitment with a friend to have lunch every week; you both move your schedules and suffer inconveniences to make that happen. But if her father fell gravely ill and she needed to miss your lunch to fly across the country to be with him, it (probably) wouldn't feel like a violation of the terms of the agreement. Apologies would not be warranted; if offered, they could be rightfully refused. In a normal friendship, it would not be the case that you had a claim on your friend's time come hell or high water; the death of a parent seems to be an unarticulated exception to the practice's norms rather than a situation in which some stronger, countervailing reason outweighed the importance of the commitment.<sup>21</sup>

---

<sup>21</sup>You can imagine a variety of reasons your friend might break her commitment and miss your lunch. If the occasion for missing the lunch were sufficiently weak (she wanted to watch a TV show), she'd be doing wrong. For some occasions (she needed to talk to her hard-to-get-a-hold-of colleague about an urgent matter), it seems that she does infringe on the claim to her time that your joint commitment gives even if what she does is on the whole permissible. She owes you an apology or other recompense (Thomson, 1990). But for the extreme case like rushing to a dying parent, no apology is warranted at all. For a compelling argument that promises in general have this feature, see Korsgaard (2009). For an application to

Julia and Charles's relationship might be like that. While all that they've encountered so far has fallen within the ambit of their extreme informational openness, this might be just the sort of case that counts as a kind of unarticulated (likely unforeseen) exception. If Julia and Charles relate to this case as we would to a friend missing lunch to be with a dying parent—i.e., as in need of no explanation or justification of the deviance from the usual—this would be a case in which the original problem dissolves: while their relationship may seem to require disclosure, it doesn't.

This phenomenon of unarticulated exceptions is widespread; most of the norms of relationships are unarticulated, both main cases and exceptions. Even when the norm is explicitly uttered (“Let's tell each other everything, ok? No secrets.”), it can still be unwarranted to expect that all the exceptions be articulated. It might go without saying, take too long, or involve unforeseen circumstances. Importantly, the very act of articulating the exceptions might weaken the psychological resolve of those adopting the norm.<sup>22</sup>

Version 2: Changing the Relationship, no reason to share

But perhaps Julia and Charles's relationship isn't like that. They might see the norms of the relationship and recommendation of the therapist as in genuine conflict; the rules really might be that strict. Relationships, however, are dynamic; they can and do change constantly. Just because the relationship's norms require disclosure when Charles and Julia begin a conversation with Julia about his therapist's recommendation doesn't mean that the relationship will require disclosure by the end of the talk. If possible, Charles should try to change

---

the communicative context, see Pallikkathayil (2017).

<sup>22</sup>See Holton (2004, 2017).

the norms of the relationship; to loosen them and create the space for the private growth he needs. If he succeeds, he would no longer be in a relationship that requires maximum disclosure; he would be in a relationship which no longer issues in that reason to share.<sup>23</sup> He could do so explicitly (“I think we should change the terms of our partnership...”) or implicitly, with his refusing to tell Julia and her accepting the change without direct comment.

Here is our second kind of reason: he should keep this information private because by doing so, he can reshape the relationship into a still better one.<sup>24</sup> This stands in addition to the fact that by doing so, he can better serve his emotional needs. Because relationships’ norms are shaped implicitly by the very things they regulate (i.e., by what we share and how we speak), it is common for us to have reasons to conceal in a manner that is in some way contrary to the relationship itself in order to make the relationship a better one. Likewise, and very frequently in a dishonest relationship, one may have reason to tell the truth not because the relationship requires it, but because by sharing, one can start to make the relationship a more honest one with better ends.

Whether this change is possible depends on how malleable the norms of their relationship are and how flexible Charles and Julia are as individuals. This is not an a priori issue; it depends on the psychology of the participants, their ability to tolerate changes in their practice, the history of the relationship and a further understanding

---

<sup>23</sup>Would this be a new relationship? Or is it the same old relationship with different norms? This question of individuation is not important for our purposes; what matters is only whether the relationship one is actually in has the given norm. Whether the original relationship changed or a new, similar relationship replaced it is irrelevant.

<sup>24</sup>Cf. MacIntyre (1994).

of why they are in this relationship to begin with. If they can change their relationship, Charles should do just that, adjusting the norms of his relationship to bring them in keeping with his evolving needs. How—explicitly, implicitly, with what words or signals—will again depend on the particulars.

### Version 3: There is Reason to Share

It may be, however, that changing the relationship requires breaking the norms of the relationship in order to change them. Imagine that at first Julia resists Charles's attempt to change the relationship; his initial attempt at keeping things private is not accepted and the relationship does not happily shift in Charles's favor. Only when Charles nevertheless breaks the norms of the relationship and insists on his privacy does Julia relent and the relationship shift. Here, we should recognize that Charles should not share for the same two reasons as above (for the immediate benefit to his emotional wellbeing and in order to better shape the relationship in the long term); but we should still recognize that he had a reason to share, generated by the relationship via transmission, but that it was overridden by his stronger reasons to conceal.

Why think Charles had an overridden reason to share at all? Why not think that this version of our case is just a counter-example, with a constitutive standard that should be disregarded? Because though defeated, the reason to share has all kinds of "trace effects." It makes sense to see Charles as infringing on a claim of Julia's, though justifiably. While he is right not to share, he still owes her an explanation; he owes her an apology; he is warranted in feeling regret, and she in

feeling resentment.<sup>25</sup> These are all hallmarks of a reason that was defeated.

What if no change is ever possible? Be it stubbornness on one or both of their parts or even a rational response to the particularities of their psychology and history together, suppose that the relationship is what it is and Charles never loses his reason to share. If it really is such a case, the question of what to do will be a matter of weighing up the reasons for and against, and depending on the particulars, things could go either way. It isn't implausible that in an imperfect world in which we build imperfect relationships with one another, we might be required to make very real sacrifices to keep the relationships we have. Charles might need to hobble his own growth if, on balance, it is less important than maintaining this relationship. But if his own growth is more important than maintaining this norm of the relationship, he might need to betray the terms of the relationship for his own sake. That doesn't necessarily mean the end of the relationship; we break the rules of our relationships all the time. It does mean, however, that as above, he has reason to feel regret, that he owes Julia recompense for the real claim she has on him that he is violating, and that she would be justified in feeling resentment for his decision. How much the relationship can tolerate this ongoing tension is again a question of the particulars and would be a large part of determining what Charles should do.

---

<sup>25</sup>See Thomson (1990).

### 3.4 Why Be Truthful? Why Discreet?

The view I've articulated at preserves what was right about the autonomy and practice arguments. The former claimed that the key to understanding the wrong of deception was the notion of autonomy, and considerations thereof are a large part of what makes for good relationships, the kind that we should create with others and which meet the requirements of transmission. The practice argument focused our attention on the centrality of trust and provided us with the basic structure of the view: our communicative practice(s) have norms we mustn't undermine. But by going in for an integrationist and local approach, we've seen far more that would otherwise go unnoticed: that the norms in question are grounded in particular relationships, that they fix a line between that which is in and out of bounds—a line whose violation in one direction amounts to dishonesty, in the other to indiscretion. Moreover, we've arrived at a unified view of communicative ethics instead of fragmented one that treats honesty and discretion separately.

By way of summary, let's first consider how all of these moving parts fit together with respect to honesty. Central to the view are the norms transmitted from the relationships we (with reason) inhabit. Though these norms permit more deception than many have thought inasmuch as they allow for the importance of concealment and privacy within a relationship, they issue in strong reasons not to deceive. Norms that permit widespread or even moderate deception normally make for frustrating, unhealthy, and unhelpful relationships in which communication is often ineffective. That is, excessively untruthful relationships should be changed or abandoned in favor of those that strike a better balance of autonomy between their participants.



In addition are all the usual consequentialist reasons not to lie. Deception normally harms, and we are often spectacularly bad at assessing those harms when we are faced with the temptation to deceive (Bok 1999; Augustine 1887b, §38).

But among the bad effects of deception is one that the relational view I've defended calls most attention to: deception's effects on relationships themselves. To deceive impermissibly isn't just to violate the terms of one's relationship; it is also to break those norms down. While consequentialists like Sidwick may have been misguided to claim that individual deceptions are wrong because they damage the norms of society by causing a (small) dissolution of social trust, deception really does have just such an effect in more local relationships. After all, relationships are dynamic—their ends and norms are shaped by us. To deceive is to kick the other out, to push her away. And if such acts are not reckoned with, they threaten to shape the rules themselves, reducing trust, increasing the distance between the participants, and normalizing what was once forbidden. People worry not just that President Trump is violating the norms that structure the relationship between president and citizen, but that he is changing that relationship altogether.

Correlatively, there is in almost any ordinary circumstance just such a reason to tell the truth: in so doing, you are building a better relationship between you and your audience, one in which they have a claim to the truth of the information you offer, and you have a reciprocal claim, a relationship governed by norms of openness and honesty, with better communication, cooperation and closeness. Yes, it will be a relationship with less autonomy reserved for each individual, but for most domains of information and with most people, it is a trade-off

we should make to some degree. And this second kind of reason to be truthful, not that the relationships' norms require it but that in so doing, one makes a relationship with better norms, is crucial.

Similarly for discretion. When we have reason to be in them, the norms of our relationships both give us permission and require us to conceal. The good relationships are those whose permissions and requirements of discretion are those that make for the flourishing of each individual in relation to one another. Like Charles, we have reason to conceal not just according to the norms of our relationships but in order to shape the relationships themselves. And to be indiscreet, likewise, can be wrong not only because it violates our relationships' norms but because it make the relationship worse off (e.g., a therapist who shares too much of his own life with his patients).

### 3.5 A Coda on Lying

On the view I've developed, no special attention has been given to lying. It falls out of the view that lying is often wrong—lying typically breaks the norms of the relationships we have reason to be in, undermines the trust within a particular relationship (thereby shifting the dynamics of the relationship in unwanted ways), not mention all the harms that consequentialists traditionally appeal to. But so, too, can withholding information, speaking truthfully but misleading (e.g., misleading by implicature), and other modes of deception. Is lying in some way special vis-a-vis deception more generally?

Many have thought that lying is in its own way worse—that over and above whatever wrongs are committed in deception is a wrong

distinctive to lying.<sup>26</sup> Abstracting the issue away from a lively and important current debate about the linguistic distinction between lying and misleading,<sup>27</sup> we can (roughly) note that that lying—as opposed to “merely misleading”—is not merely a misuse of implicature but also of assertion (or something along those lines).<sup>28</sup> Put in terms of a globalist practice argument, we might consider two important communicative practices: implicature and assertion. Lying runs afoul of both, implicature only of one; it’s worse.

Plenty of cases support the conclusion that lying is worse than misleading. Imagine Claudia conceals her cancer from a friend; we can easily imagine their relationship such that her friend would feel a difference in the wrong done if Claudia lied or merely misled. We’ve also seen cases in which significant misleading is permitted while (almost all) lying is forbidden: the courtroom.

But others resist this conclusion and insist that there is no important distinction between lying and misleading. They point to cases in which the misleading and lying seem to be equally wrong, e.g., when we considered Claudia misleading her romantic partner about her health;<sup>29</sup> lying or mere misleading—it doesn’t matter. And complicating the whole question are lingering issues about the do/allow distinction, however it is to be drawn, as lies are typically (if not

---

<sup>26</sup>Ross (1930, 22-28); Chisholm and Feehan (1977); Adler (1997); Strudler (2010); Webber (2013). For the view that lying and deceiving are morally equivalent see Saul (2012b,a), and for the surprising, minority view that non-lying deception is worse, see Rees (2014).

<sup>27</sup>See, for instance, Carson (2006, 2010); Fallis (2010); Saul (2012b). For an outstanding overview, see Mahon (2015).

<sup>28</sup>As in Chisholm and Feehan (1977).

<sup>29</sup>Motivated by cases like these, Saul argues that lying and misleading are morally on a par (2012b). In Saul’s case, someone deceives a would-be sexual partner about having an STI; lying and misleading, she says, are equally bad. She’s right, in this case, but not in general. For a nuanced view on the topic, see (Pallikkathayil, 2017, pp. 14-16).

always) doings but other forms of deception can be mere allowings.

A relationships approach can help us make sense of all this. Both sides of this dispute are right... but only about their favored cases. Whether lying is in fact worse depends on the constitutive norms of the relationship in question and whether, given its ends, drawing a normative distinction between lying and other modes of deception makes sense. In court, the ends of the relational practice of jury, advocate and judge are best served by drawing a very sharp distinction between (disallowed) lying (i.e., perjury) and permissible misleading (i.e., lawyering). In a relationship of the deepest trust, like a healthy romantic partnership, that distinction often serves no valuable purpose. What matters is trust and openness; what need there is for privacy and concealment can and should be achieved by other means, e.g., saying "I'd rather not go there. I need you to trust me on this." (A move that wouldn't work in court....) The trust of an intimate relationship allows that conversational move to be made successfully.

But most relationships, most of the time, will recognize a difference between lying, merely misleading and other modes of deception. In order to recognize varying degrees of communicative infringements, it makes sense that they do so. We can consider a spectrum of ever-more-deceptive communicative moves<sup>30</sup>, running roughly from staying silent to dodging the issue to misleading to lying, with each step further along being more communicatively forceful. What I mean here is that the things further along leave less room for the listener to respond. If I stay silent, all possibilities are left open to you; if I dodge the question, I have in some way limited the conversation. You can still

---

<sup>30</sup>For a well-developed view of this, see [redacted], which sees a similar spectrum as being defined by speech act's being more or less informative cached out in terms of Grice's Maxims of Quantity.

push me, but that takes a kind of work on your part; you have to resist my move. The resistance required is still greater if I mislead without lying. Again, you can pin me down (“I know you said you didn’t drive the car, but that still doesn’t tell me whether you scratched it!”), but the burdens are greater. A lie requires the most work to resist; it requires not only pinning down but retraction. Put another way, as we move along the spectrum, the force we exert on another’s doxastic agency is increased. And in general, the more forceful the move, the more it impinges upon the other’s autonomy. (I say “in general” because this is content-dependent. Claudia’s lying about the color of the car down the street would have a much lesser impact on Paige’s agency than keeping silent about her cancer.)

By having the strength of the reasons not to conceal increase along the spectrum, a relationship can make more normative space for concealment, an arrangement that leaves more autonomy in the hands of individuals. So if for Claudia and her friend, lying is in fact worse than merely misleading, it could be that the one is permissible where the other isn’t. I do not mean that one is permitted by the norms of the relationship (as in the case of the coworker); I mean permitted given countervailing reasons to deceive. By making the one worse, the relationship makes there be a difference in the strength of reasons. Thus, it might matter enough to Claudia that the information be kept private to justify misleading but not so much that lying is justified.<sup>31</sup>

---

<sup>31</sup>The phenomenon of non-deceptive lies neatly falls out of this picture. Imagine a witness at a mob trial who, lying to protect himself, claims he did not see the defendant at the scene of the crime. He does not intend the jury to believe him—he hopes they see through his thinly veiled lie—he just has to say what he needs to in order to appease the mob (Carson, 2006). Such lies are at once continuous with and distinct from deceptive lies. They break communicative norms of a relationship, but do so in a way that is intentionally ineffective. They

I'm not suggesting that ordinary people have had thoughts like these or structured their relationships because of these considerations. But I do think that this kind of functional explanation for the distinction can explain the forces that shape relationships' communicative norms. And note that these norms can be context, purpose and topic sensitive; merely misleading someone about having an STI in order to have sex is just as bad as lying (Saul, 2012b), even absent a relationship that more generally elides the distinction; the norms governing even a relatively non-intimate sexual relationship will still be sensitive to purpose, potential harms and the importance of the information.

It is for these reasons that contrary to the suggestions of some, we need not settle on a definition of lying before doing any normative investigation into it.<sup>32</sup> The definition of lying, i.e., the boundaries between lying on the one hand, and other modes of deception on the other, depends on the normative particularities of our situation. Here, as elsewhere, the truth about lying, deception and how we should speak to one another is, for better and for worse, very much in our hands.

---

refuse to share what is to be shared, but they do not attempt to deceive the listener about the truth of what is to be shared. Depending on the case, they may also be more rude than merely refusing to answer a question inasmuch as they more brazenly flout the norms of the relationship.

<sup>32</sup>As in Carson (2006); Fallis (2010). See also Pallikkathayil (2017), who likewise sees the normative analysis as independent of the definitional one.

# Bibliography

- Adler, J. E. (1997). Lying , Deceiving , or Falsely Implicating. *The Journal of Philosophy*, 94(9):435–452.
- Arpaly, N. and Schroeder, T. (2013). *In Praise of Desire*. Oxford University Press, Oxford.
- Arpaly, N. and Schroeder, T. (2015). A Causal Theory of Acting for Reasons. *American Philosophical Quarterly*, 52(2).
- Augustine (1887a). *Enchiridion*. In Schaff, P., editor, *Nicene and Post-Nicene Fathers of the Christian Church*, volume 3. Christian Literature Publishing Co., Buffalo, NY.
- Augustine (1887b). *On Lying*. In Schaff, P., editor, *Nicene and Post-Nicene Fathers of the Christian Church*. Christian Literature Publishing Co., Buffalo, NY.
- Bacon, F. (2009). *The Essays or Counsels, Civil and Moral*. Project Guttenberg, Urbana, Illinois.
- Barlow, J. and Nadeau, J.-B. (2016). *The Bonjour effect : the secret codes of French conversation revealed*.
- Bentham, J. (1907). *An Introduction to the Principles of Morals and Legislation*. Clarendon Press, Oxford.
- Bok, S. (1989). *Secrets: On the Ethics of Concealment and Revelation*. Vintage Books, New York.
- Bok, S. (1999). *Lying: Moral Choice in Public and Private Life*. Vintage Books, New York, 2nd edition.
- Bonhoeffer, D. (1965). What is meant by ‘telling the truth’? In Bethge, E., editor, *Ethics*, pages 363–72. Macmillan Company, New York.

- Boolos, G. (1984). To Be Is To Be a Valuable of a Variable (or to Be Some Values of Some Variables). *Journal of Philosophy*, 81(8):430–50.
- Bratman, M. (1999a). Intention, plans, and practical reason. Center for the Study of Language and Information.
- Bratman, M. E. (1999b). Identification, Decision, and Treating as a Reason. In *Faces of Intention: Selected Essays on Intention and Agency*, pages 185–206. Cambridge University Press, Cambridge.
- Brown, C. (2011). Consequentialize This. *Ethics*, 121(4):749–771.
- Butler, B. J. (1740). Dissertation of the Nature of Virtue. In Raphael, D. D., editor, *British Moralists*, pages 378–386. Hackett Publishing Company, Indianapolis, 3rd edition.
- Carson, T. L. (2006). The Definition of Lying. *Nous*, 40(2):284–306.
- Carson, T. L. (2010). *Lying and Deception: Theory and Practice*. Oxford University Press, Oxford.
- Chang, R. (1997). Incommensurability, incomparability, and practical reason. Harvard University Press, Cambridge.
- Chang, R. (2002). The Possibility of Parity. *Ethics*, 112:659–688.
- Chang, R. (2013). Grounding practical normativity: going hybrid. *Philosophical Studies*, 164(1):163–187.
- Chisholm, R. M. and Feehan, T. D. (1977). The Intent to Deceive. *The Journal of Philosophy*, 74(3):143–159.
- Cohen, G. A. (2011). Rescuing Conservatism. In *Reasons and Recognition: Essays on the Philosophy of T.M. Scanlon*, pages 203–226. Oxford University Press, Oxford.
- Confucius (1997). *The Analects of Confucius*. W.W. Norton, New York.
- Cottingham, J. (1983). Ethics and Impartiality. *Philosophical Studies*, 43(1):83–99.
- Cumberland, R. (1727). *A Treatise of the Laws of Nature*. Liberty Fund, Indianapolis.
- Curtis, C. P. (1951). The Ethics of Advocacy. *Stanford Law Review*, 4(1):3–23.



- D'Arms, J. and Jacobson, D. (2000a). Sentiment and Value. *Ethics*, 110(4):722–748.
- D'Arms, J. and Jacobson, D. (2000b). The Moralistic Fallacy : On the ' Appropriateness ' of Emotions. *Philosophy and Phenomenological Research*, 61(1):65–90.
- Darwall, S. (2006). *The Second-Person Standpoint: Morality, Respect, and Accountability*. Harvard University Press, Cambridge, MA.
- Darwall, S. (2013). Morality's Distinctiveness. In *Morality, Authority and the Law: Essays in Second-Personal Ethics I*, pages 3–19. Oxford University Press, Oxford.
- Darwall, S. (2018). Love's Second-Personal Character: Reciprocal Holding, Beholding and Upholding. In Kroeker, E. E. and Schaubroeck, K., editors, *Love , Reason and Morality*, chapter 2016. Routledge, New York.
- Davies v. Coverstone (1952). 38 Cal.2d 315, 323 [239 P.2d 876].
- Dreier, J. (2011). In defense of consequentializing. In Timmons, M., editor, *Oxford Studies in Normative Ethics*, volume 1. Oxford University Press, Oxford.
- Dynel, M. (2011). A Web of Deceit: A Neo-Gricean View on Types of Verbal Deception. *International Review of Pragmatics*, 3:139–167.
- Ebels-Duggan, K. (2009). Against Beneficence: A Normative Account of Love. *Ethics*, 119(1):142–170.
- Evans, D. (1969). *The Logic of Self Involvement*. Herder & Herder, New York.
- Fallis, D. (2010). Lying and Deception. *Philosopher's Imprint*, 10(11):1–22.
- Faulkner, P. (2007). What Is Wrong with Lying? *Philosophy and Phenomenological Research*, 75(3):535–557.
- Foot, P. (2001). *Natural goodness*. Clarendon.
- Frankfurt, H. G. (2004). *The Reasons of Love*. Princeton University Press, Princeton.
- Freud, S. (2010). *Civilization and Its Discontents*. Norton, New York.

- Gandhi, M. (1957). *Gandhi: An autobiography — The Story of My Experiments With Truth*. Beacon Press, Boston.
- Gill v. Hearst Publishing Co. (1953). 40 Cal.2d 224 [253 P.2d 441].
- Goodman, L. E. (2008). *Love Thy Neighbor As Thyself*. Oxford University Press, Oxford.
- Griffiths, P. J. (2004). *Lying: An Augustinian Theology of Duplicity*. Brazos Press, Grand Rapids.
- Grotius, H. (2012). *On the Law of War and Peace*. Cambridge University Press, Cambridge.
- Hare, C. (2010). Take the sugar. *Analysis*, 70:237–247.
- Hedden, B. (2015a). Options and Diachronic Tragedy. *Philosophy and Phenomenological Research*, 90(2):423–451.
- Hedden, B. (2015b). *Reasons without Persons : Rationality, Identity, and Time*. Oxford University Press, Oxford.
- Hedden, B. (2015c). Time-slice rationality. *Mind*, 124(494):449–491.
- Hobbes, T. (1994). *Leviathan*. Hackett Publishing Company, Indianapolis, Indiana.
- Holton, R. (2004). Rational Resolve. *Philosophical Review*, 113(4):507–535.
- Holton, R. (2017). We Don't Torture : Moral Resolutions , Temptation , and the Doctrine of Double Effect. *Journal of the British Academy*, 5:309–329.
- Hume, D. (2007). *A Treatise of Human Nature*. Oxford University Press, Oxford.
- Hurka, T. (2014). *British Ethical Theorists from Sidwick to Ewing*. Oxford University Press, Oxford.
- Hutcheson, F. (1738). *An Inquiry into the Original of Our Ideas of Beauty and Virtue in Two Treatises*. In Raphael, D. D., editor, *British Moralists*. Hackett Publishing Company, Indianapolis, 4th edition.
- Hutcheson, F. (1755). *A System of Moral Philosophy*. A. Millar, Digitized by Univ. of California Libraries, London.

- Irfa, U. (2018). Why soccer players take dives. *Vox*.
- Jeske, D. (2008). *Rationality and Moral Theory*. Routledge, New York.
- Johnston, M. (2010). *Surviving death*. Princeton University Press, Princeton.
- Jollimore, T. (2011). *Love's Vision*. Princeton University Press, Princeton.
- Jollimore, T. (2017). Love as "Something in Between". In Grau, C. and Smuts, A., editors, *The Oxford Handbook of Philosophy of Love*. Oxford University Press, Oxford.
- Jollimore, T. (2019). *Constancy and Requirement in Romantic Love*. Unpublished Manuscript.
- Kagan, S. (1991). *The Limits of Morality*. Oxford University Press.
- Kagan, S. (2000). Evaluative Focal Points. *Morality, Rules, and Consequences*, pages 134–155.
- Kant, I. (1993). On a Supposed Right to Lie because of Philanthropic Concerns. In *Grounding for the Metaphysics of Morals, with On a Supposed Right to Lie because of Philanthropic Concerns*. Hackett Publishing Company, Indianapolis.
- Kant, I. (1998a). *Groundwork of the Metaphysics of Morals*. Cambridge University Press, Cambridge.
- Kant, I. (1998b). *Religion within the Boundaries of Mere Reason*. Cambridge University Press, Cambridge.
- Kant, I. (2001). *Lectures on Ethics*. Cambridge University Press, Cambridge.
- Kant, I. (2017). *The Metaphysics of Morals*. Cambridge University Press, Cambridge, 2nd edition.
- Keller, S. (2013). *Partiality*. Princeton University Press, Princeton.
- Kierkegaard, S. (1995). *Works of Love*. Princeton University Press, Princeton.
- Kolodny, N. (2003). Love as Valuing a Relationship. *Philosophical Review*, 112(2):135–189.

- Kolodny, N. (2010a). Which Relationships Justify Partiality? General Consideration and Problem Cases. In Feltham, Brian; Cottingham, J., editor, *Partiality and Impartiality: Morality, Special Relationships, and the Wider World*, pages 169–193. Oxford University Press, Oxford.
- Kolodny, N. (2010b). Which Relationships Justify Partiality? The Case of Parents and Children. *Philosophy and Public Affairs*, 38(1):37–75.
- Korsgaard, C. M. (1986). The Right to Lie: Kant on Dealing with Evil. *Philosophy Public Affairs*, 15(4):325–349.
- Korsgaard, C. M. (1996). *The Sources of Normativity*. Cambridge University Press, Cambridge.
- Korsgaard, C. M. (1997). The Normativity of Instrumental Reason. In Cullity, G. and Gaut, B., editors, *Ethics and Practical Reason*. Clarendon Press, Oxford.
- Korsgaard, C. M. (1999). The General Point of View: Love and Moral Approval in Hume’s Ethics. *Hume Studies*, XXV(1):3–41.
- Korsgaard, C. M. (2009). *Self-Constitution : Agency, Identity, and Integrity*. Oxford University Press, Oxford.
- Leibniz, W. G. (1969). *Discourse on Metaphysics*. In Loemker, L. E., editor, *Philosophical Papers and Letters*. University of Chicago Press, Chicago.
- Lewis, D. (1983). *Language and Languages*. In *Philosophical Papers*, volume 2, pages 3–35. Oxford University Press, Oxford.
- MacIntyre, A. (1994). Truthfulness, Lies, and Moral Philosophers: What Can We Learn from Mill and Kant? *The Tanner Lectures on Human Values*, 16:307–361.
- Mahajan, K. (2016). *My Struggle with American Small Talk*. The New Yorker.
- Mahon, J. E. (2003). Kant on Lies, Candour and Reticence. *Kantian Review*, 7:101–133.
- Mahon, J. E. (2009). The Truth About Kant On Lies. In Martin, C., editor, *The Philosophy of Deception*. Oxford University Press.

- Mahon, J. E. (2015). The Definition of Lying and Deception. The Stanford Encyclopedia of Philosophy.
- Mandela, N. (2018). The Prison Letters of Nelson Mandela. Liveright Publishing Corporation, New York.
- Melvin v. Reid (1931). 112 Cal. App. 285, 290 [297 P. 91].
- Mill, J. S. (2001). Utilitarianism. Hackett Publishing Company, Indianapolis, 2nd edition.
- More, H. (1666). An Account of Virtue, or, Dr. Henry More's Abridgement of Morals [Enchiridion Ethicum]. London (access via Henry E. Huntington Library and Art Gallery & Early English Books Online).
- Muñoz, D. (2019). The Paradox of Consequentialism. Unpublished Manuscript.
- Murdoch, I. (1970). The sovereignty of good. Routledge, London.
- Murphy, L. (2000). Moral Demands in Nonideal Theory. Oxford University Press, Oxford.
- Nagel, T. (1998). Concealment and Exposure. *Philosophy & Public Affairs*, 27(1):3–30.
- Niebuhr, R. (1935). An Interpretation of Christian Ethics. Harper & Brothers, New York.
- Niebuhr, R. (1957). Love and Justice: Selections from the Shorter Writings of Reinhold Niebuhr. Westminster/John Knox Press, Louisville.
- Nozick, R. (1989). The Examined Life. Simon & Schuster, New York.
- Nygren, A. (1953). Agape and Eros. SPCK, London.
- Orwell, G. (1949). Reflections on Gandhi. In Carey, J., editor, *Essays*, pages 1349–56. Alfred A. Knopf, New York.
- Outka, G. (1972). Agape: An Ethical Analysis. Yale University Press, New Haven.
- Pallikkathayil, J. (2010). Deriving Morality from Politics: Rethinking the Formula of Humanity. *Ethics*, 121(1):116–147.

- Pallikkathayil, J. (2017). The Truth About Deception. *Philosophy and Phenomenological Research*, pages 1–20.
- Parfit, D. (2011). *On What Matters*. Oxford University Press, Oxford.
- Perry, R. B. (1909). *The Moral Economy*. Charles Scribner's Sons, New York.
- Pettit, P. (1997). Love and Its Place in Moral Discourse. In Lamb, R., editor, *Love Analyzed*, pages 153–163. Westview Press, Boulder, CO.
- Porten v. University of San Francisco (1976). 64 Cal. App. 3d 825, 828 [134 Cal. Rptr. 839].
- Rawls, J. (1995). Two Concepts of Rules. *The Philosophical Review*, 64(1):3–32.
- Rees, C. F. (2014). Better lie! *Analysis*, 74(1):59–64.
- Ripstein, A. (2009). *Force and Freedom: Kant's Legal and Political Philosophy*. Harvard University Press.
- Ross, W. D. (1930). *The Right and The Good*. Clarendon Press, Oxford.
- Ross, W. D. (1939). *Foundations of Ethics: The Gifford Lectures Delivered in the University of Aberdeen*. Clarendon Press, Oxford.
- Samuel v. Curtis Pub. Co. (1954). (N.D. Cal. 1954) 122 F. Supp. 327.
- Saul, J. (2012a). Just go ahead and lie. *Analysis*2, 72(1):3–9.
- Saul, J. M. (2012b). *Lying, Misleading, and What is Said: An Exploration in Philosophy of Language and in Ethics*. Oxford University Press, Oxford.
- Scheffler, S. (1997). Relationships and Responsibilities. *Philosophy and Public Affairs*, 26(3):189–209.
- Scheffler, S. (2010). Morality and Reasonable Partiality. In Feltham, B. and Cottingham, J., editors, *Partiality and Impartiality: Morality, Special Relationships, and the Wider World*, pages 98–130.
- Schneewind, J. B. J. B. (1998). *The Invention of Autonomy : A History of Modern Moral Philosophy*. Cambridge University Press, Cambridge.

- Setiya, K. (2010). *Reasons Without Rationalism*. Princeton University Press.
- Setiya, K. (2011). Reasons and causes. *European Journal of Philosophy*, 19(1):129–157.
- Setiya, K. (2014). Love and the Value of a Life. *Philosophical Review*, 123(3):251–280.
- Setiya, K. (2019). *Other People*. Unpublished Manuscript2.
- Shiffrin, S. V. (2014). Lies and the Murderer Next Door. In *Speech Matters: On Lying, Morality and the Law*, chapter 1, pages 5–46. Princeton University Press, Princeton.
- Sidgwick, H. (1981). *The Methods of Ethics*. Hackett Publishing Company, Indianapolis, 7th edition.
- Singer, P. (1972). Famine, affluence, and morality. *Philosophy & Public Affairs*, pages 229–243.
- Smith, J. (1660). *Select Discourses*. London (Access via Google Books), 3rd edition.
- Spaemann, R. (2012). *Love and the Dignity of Human Life: On Nature and Natural Law*. W.B. Eerdmans Publishing Co.
- Stenius, E. (1967). Mood and Language-Game. *Synthese*, 17(3):254–74.
- Stocker, M. (1976). The Schizophrenia of Modern Moral Theories. *Journal of Philosophy*, 73(14):453–466.
- Stocker, M. (1981). Values and Purposes: The Limits of Teleology and the Ends of Friendship. *The Journal of Philosophy*, 78(12):747–765.
- Strudler, A. (2010). The Distinctive Wrong in Lying. *Ethical Theory and Moral Practice*, 13(2):171–79.
- The Mahabharata (2016). *The Mahabharata*. In DeLapp, K. and Henkel, J., editors, *Lying and Truthfulness*, chapter 4,7, pages 63–69,100–106. Hackett Publishing Company, Indianapolis.
- Thomas Aquinas (2016). *Summa Theologiae*. In DeLapp, K. and Henkel, J., editors, *Lying and Truthfulness*, chapter 11, pages 158–184. Hackett Publishing Company, Indianapolis.

- Thompson, M. (2008). *Life and Action*. Harvard University Press.
- Thomson, J. J. (1990). *The realm of rights*. Harvard University Press, Cambridge.
- Timperley v. Chase Collection Service (1969). 272 Cal. App. 2d 697, 700 [77 Cal. Rptr. 782].
- Unger, P. K. (1996). *Living high and letting die : our illusion of innocence*. Oxford University Press, Oxford.
- Velleman, J. D. (1999). Love as a Moral Emotion. *Ethics*, 109(2):338–374.
- Way, J. (2012). Transmission and the Wrong Kind of Reason. *Ethics*, 122(3):489–515.
- Webber, J. (2013). Liar! *Analysis*, 73(4):651–59.
- Wedgwood, R. (2006). The normative force of reasoning. *Noûs*, 40(4):660–686.
- Williams, B. A. O. (1981). Persons, Character and Morality. In *Moral Luck: Philosophical Papers, 1973-1980*, pages 1–19. Cambridge University Press, Cambridge.
- Williams, B. A. O. (2002). *Truth and Truthfulness*. Princeton University Press, Princeton.
- Wolf, S. (1982). Moral Saints. *The Journal of Philosophy*, 79(8):419–439.
- Wolf, S. R. (2010). *Meaning in life and why it matters*. Princeton University Press, Princeton.
- Wolterstorff, N. (2011). *Justice in Love*. Wm. B. Eerdmans Publishing Co., Grand Rapids, Michigan.
- Zembya, J. S. (1993). Aristotle on Lying. *Journal of the History of Philosophy*, 31(1):7–29.