



Loving, Valuing, Regretting, and Being Oneself

Citation

Na'aman, Oded. 2015. Loving, Valuing, Regretting, and Being Oneself. Doctoral dissertation, Harvard University, Graduate School of Arts & Sciences.

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:23845400>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Loving, Valuing, Regretting, and Being Oneself

A dissertation presented

by

Oded Naaman

to

The Department of Philosophy

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Philosophy

Harvard University

Cambridge, Massachusetts

June 2015

© 2015 Oded Naaman

All rights reserved.

Loving, Valuing, Regretting, and Being Oneself

Abstract

A meaningful life involves loving people and valuing things. We typically love our spouses, parents, children, siblings, and friends, and value our projects, activities, causes, and ideals. In virtue of such attachments, a meaningful life is also susceptible to regret of a distinctively personal kind. Our regrets about the misfortunes and harms that befall the people we love and the things we value reflect the extent to which we are implicated in the fate of those people and things, the extent, that is, to which our attachments determine who we are. What are the reasons to which we respond in loving a particular person or valuing a particular thing, and how do these reasons explain personal regret as well as our conception of who we are? This is the question this dissertation aims to address.

The primary thesis of the first chapter is that we have reasons to love particular individuals as such. The primary thesis of the second chapter is that different individuals rationally value different objects in different ways, while one and the same object can be rationally valued very much by some and not at all by others. The primary thesis of the third chapter is that our present attachments give us reason to regret or affirm the past.

All three chapters make use of the distinction between reasons *of* attachment and reasons *for* attachment. Reasons *of* attachment are the reasons a person takes him- or herself to have with regard to the individual or object he or she loves or values, and reasons *for* attachment justify or warrant a person's attachments as well as enable the

person's reasons *of* attachment. For example, the fact that a person is kind might be a reason *for* being his or her friend, but the fact that the person is free this afternoon might be a reason *of* friendship to spend this afternoon with him or her. Alternatively, the fact that Ness-Ziona is my hometown may be a reason *for* valuing it, but the fact that the orchards of my childhood have been replaced by suburban neighborhoods is a reason *of* valuing to lament the change Ness-Ziona has undergone.

In addition to the three, positive theses, all three chapters discuss ways in which our reasons *for* and *against* attachment might clash with our reasons *of* attachment with regard to a particular person or object. While reasons *for* attachment and reasons *of* attachment often complement each other, they might also pull in opposite directions. When we have most reason to let go of the person we love, of the place that is our home, or of the vocation we cherish, and letting go means failing or abandoning this person, place, or vocation, then the two kinds of reasons may set us against ourselves, as it were. Thus, this dissertation aims to shed light on the rationality of a meaningful life, but it also accounts for inevitable crises of meaning, without which a meaningful life would not be such a tremendous achievement.

Table of Contents

Acknowledgments	vi
Prologue and Introduction	1
Chapter I: Loving	11
<i>Introduction</i>	11
1. A Puzzle about Partiality	13
2. A Puzzle about Particularity	19
3. A Puzzle about Necessity	35
Chapter II: Valuing	48
<i>Introduction</i>	48
1. What Is Valuing?	51
2. Valuing and Believing Valuable	58
3. What Makes Valuing Personal?	69
4. A Lingering Puzzle about Devotion	75
Chapter III: Regretting and Being Oneself	90
<i>Introduction</i>	90
1. Not Regretting	93
2. Regretting a Mistake	104
3. Regretting a Loss	132
4. Regretting an Attachment	142

Acknowledgments

Most of the time, writing this dissertation seemed like an attempt to find my way in a large, dark chamber, with only the voices of my advisors to guide me. Tim Scanlon, Christine Korsgaard, Dick Moran, and Selim Berker have been patient, attentive, and endlessly illuminating. I have been astounded and humbled by their philosophical acumen and sensibilities, by the breadth of their philosophical world-views, and by the depth of their philosophical understanding. I am profoundly grateful for the opportunity to read them, listen to them, read them again, argue with them, learn from them, concede their points, read them once more, and be understood by them. Their voices will continue to echo and counsel me in the dark chambers that lie ahead.

I consider myself lucky to have spent these years studying with remarkable people as my fellow graduate students. For their camaraderie, insight, laughter, despair, and inexorable passion for philosophy, I thank: Diana Acosta Navas, Olivia Bailey, James Bondarchuk, Byron Davies, Tom Fehse, Jeremy Fix, Johann Frick, Marc Gasser, Micha Glaeser, Paul Julian, Douglas Kremm, Enoch Lambert, Dave Langois, Céline Leboeuf, Patricia Marechal, Emily McWilliams, Elizabeth Miller, Alex Prescott-Couch, James Reich, Ronni Sadovsky, Paul Schofield, Zeynep Soysal, Aleksy Tarasenko-Struc, Kate Vredenburgh, and Bernardo Zacka.

The Edmond J. Safra Center for Ethics at Harvard has graciously supported my research in the academic year of 2013-2014. I wish to thank Eric Beerbohm, who led the Graduate Seminar at the Safra Center for Ethics, as well as the other participants in the seminar: Joelle Abi-Rached, Ryan Davis, Mark Hanin, Philipa Hetherington, Tae-Yeoun

Keum, Charles Lesch, and Aline-Florence Manent.

The Harvard Philosophy and Literature Reading Group, which I attended for three wonderful years, had a lasting impact on me. It would not be an exaggeration to say that this reading group shaped who I am today. I cannot imagine my time at Harvard without it, and would like to thank the organizers of this group: Philip Fisher, Dick Moran, and Elaine Scarry, as well as the many other participants. I would like to add a special note of gratitude for Elaine Scarry, for her unyielding faith in me, for her advice and encouragement.

Finally, I would like to thank my Mother, Yael Naaman, my Father, Ron Naaman, and my three sisters: Einat Sitbon, Adi Diner, and Irit Naaman. They have always believed in me more than I was able to; my love for them is endless. And to thank, also, Yaara Schattner, who knows who she is, and who we are together.

This dissertation is dedicated to Yehoshua Kenaz, a beacon of truthfulness.

Prologue and Introduction

Keeping Things Whole/Mark Strand

In a field
I am the absence
of field.
This is
always the case.
Wherever I am
I am what is missing.

When I walk
I part the air
and always
the air moves in
to fill the spaces
where my body's been.

We all have reasons
for moving.
I move
to keep things whole.

(Strand 2014, 78)

After Mark Strand's passing in November 2014, the *New York Review of Books* published a short speech about his poetry that was originally delivered by Joseph Brodsky 28 years earlier. Brodsky delivered the speech as an introduction to a reading by Strand at the American Academy of Poets in New York City on November 4, 1986. In his opening remarks, Brodsky describes his attitude toward Strand's poems.

It's a tall order to introduce Mark Strand because it requires estrangement from what I like very much, from something to which I owe many moments of almost physical happiness—or to its mental equivalent. I am talking about his poems—as well as about his prose, but poems first.

A man is, after all, what he loves. But one always feels cornered when asked to explain why one loves this or that person, and what for. In order to explain it—which inevitably amounts to explaining oneself—one has to try to love the object of one's attention a little bit less. I don't think I am capable of this feat of objectivity, nor am I willing even to try. In short, I feel biased about Mark Strand's poems, and judging by the way his work progresses, I expect to stay biased to the end of my days. (Brodsky 1986/2015)

Brodsky suggests that while one may have reasons to love—reasons that justify, warrant, and make sense of one’s love—recognizing and articulating these reasons may require a certain distance from one’s love and therefore a certain distance from oneself. And although Brodsky claims to be unable to distance himself in this way, and says that he is not willing even to try, he does go on to offer a very eloquent explanation of his reasons for loving Strand’s poetry. Indeed, even in the last sentence of the passage quoted above, and immediately after rejecting the demand to justify his love, Brodsky predicts, on the basis of the progress of Strand’s work, that he will continue to have reason to love his poetry.

Perhaps at the end of his speech, having articulated and grasped his reasons for loving Strand’s poetry, Brodsky loved it even more and was closer to it than ever before. Reflection on one’s reasons to love may require distance and momentary disengagement from the object of one’s love, for reflection on one’s love involves questioning one’s attachment to the object of love. But while such questioning may eventually undermine one’s love, it may also lead one back to it with greater fervor. If, on the other hand, one fails ever to consider the reasons for one’s love, then one runs the risk of being blinded by love to the point of misunderstanding its boundaries or misconceiving its object. Loving a person or a thing involves moving one’s attention back and forth, from one’s object of love to one’s reasons to love, and it requires doing so appropriately, that is, in accordance with one’s reasons. While there are tensions between these points of view—one internal, the other external to our love—we cannot forgo either one. In love, too, we must move to keep things whole.

*

This dissertation is about the relation between the persons and things that are of fundamental importance in an individual's life and the normative reasons for which they are important. Particular individuals, relationships, activities, projects, places, causes, ideals, cultures, and traditions are fundamentally important to us. We direct most of our attention and concern to them, and we normally say that they give meaning to our lives or that they make our lives worthwhile. They are the things we are most devoted to and whose absence, destruction, failure, or loss we most regret and mourn. We are attached to them.

And yet we may reflect on our attachments to such things, justify our attitudes toward them, as well as defend the actions we take for their sake or with regard to them. In other words, there are normative reasons for and against the importance of such things in our lives; there are reasons for and against our particular attachments. It might therefore seem that what truly is of most fundamental importance to us are not those particular persons, activities, projects, causes, ideals, cultures, and traditions, to which we are attached, but the facts that give us reasons for holding them in such high regard.

For example, it might seem that what is most important to me is not my close friend, but the facts that give me reason for being her friend; or that what is most important to her is not her tradition, but the facts that give her reason for endorsing it; or that what is most important to him is not his project, but the facts that give him reason to pursue it. After all, if it were not for the facts that give us reason for attachment, our attachments would not have been worthwhile or justified. Thus the fact that we may

acknowledge our reasons for having certain attachments might seem to suggest that the facts that give us reasons for attachment are more fundamentally important to us than the objects to which we are attached.

In light of these observations, we may feel “cornered,” as Brodsky puts it—compelled to choose between a view on which the objects of our attachments are of fundamental importance to us and a view on which there are reasons for and against our attachments. The former view may seem to invoke a Humean picture of normativity, according to which reason is the slave of attachments, as it were. On such a view, there can be no rational basis for our attachments, which are the bedrock of rational justification. The latter view, by contrast, seems to adopt a broadly rationalist approach, according to which reason is the charioteer and attachments are his winged horses. On such a view, the importance of the objects of our attachments can only be secondary to the importance of the facts that give us reason to be attached to them.

The Humean view seems to imply that we cannot truly endorse, disown, adopt or let go of our own attachments on the basis of judgments about reasons for and against them. The rationalist view, on the other hand, might seem to preclude disproportionate devotion to one object of attachment at the expense of another, equally worthy object. And when an object of attachment is lost, reason might seem to prescribe a worthy substitution and rule out the possibility of irredeemable loss.

These upshots are unacceptable. The fact that we reflect on the normative status of our attachments, that we justify, criticize, endorse, or disown attachments, suggests that we are committed to there being reasons that count for and against attachments. But

the fact that we care most about the objects of our attachments, and that devotion and the possibility of loss are essential to the phenomenon of attachment, suggests that the objects of our attachments, and not the facts that give us reason for attachment, are of fundamental importance to us.

This dissertation aims to do justice to both observations by rejecting the apparent dichotomy between the idea that there are reasons for attachment and the idea that objects of attachment are of fundamental importance. The Humean view and the rationalist view sketched above are caricatures that should be dispensed with. The facts that give us reason for attachment to a particular person, activity, or goal, are not the facts to which we are attached, they are not the focus of our concern, and we need not act for their sake. Thus, the first and primary claim to be promoted and examined here is that our reasons to love and value certain particular persons and things are reasons to attribute fundamental importance to them, reasons to devote ourselves to them as the particulars they are, and reasons to regret or mourn their irredeemable loss.

The second major idea of this dissertation is that there are normally genuine conflicts between, on the one hand, our reasons for and against attachment and, on the other hand, our commitments to the objects of our attachment. At the very least, being properly attached to a person or thing precludes constantly considering the justification for one's attachment. But commitment to what one values or loves might require much more than lack of self-doubt. Loving a person may often involve believing in him or her, and in the relationship we share with him or her, even when it seems to us that we have reason to forgo our love and withdraw from the relationship. Similarly, valuing a certain

goal, project, tradition, or culture may often involve pursuing it or engaging with it even when it seems to us that we have good reason to give it up or disown it. And even when we should indeed let go, withdraw from, or terminate an attachment, in doing so we may be failing our commitment to the object of attachment. This would be a warranted failure, but a failure nonetheless.

Although we can have reasons to ascribe fundamental importance to the objects of our attachments, and such reasons partly explain the genuine importance of these objects, the objects of our attachments normally command devotion and commitment that withstand strong reasons against the continuation of our attachment. By rejecting the strict dichotomy between reasons for attachment and the fundamental importance of objects of attachment, we may come to see the different ways in which the two sets of normative considerations—external and internal to our attachments—are mutually reinforcing, as well as the ways in which they come into conflict. Thus, this dissertation aims to shed light on the rationality of a meaningful life, but it also accounts for inevitable crises of meaning without which a meaningful life would not be such a tremendous achievement.

*

We begin with love. In the first chapter I consider three puzzles about the idea that there are reasons for loving a particular person: a puzzle about love's partiality, a puzzle about love's particularity, and a puzzle about love's necessity. The puzzle of partiality concerns our reasons for ascribing greater significance and attention to our beloveds than we ascribe to other individuals who are no less deserving of it. The puzzle

of particularity concerns our reasons for loving particular persons as such given that our love is justified on the basis of our beloveds' properties. The puzzle of necessity concerns the rational basis of our commitment to beloveds that we are not rationally required to love. I argue that we may disarm these puzzles by drawing a distinction between reasons *to* love and reasons *of* love. While reasons *to* love warrant, justify, or make sense of our love for a person, as well as rationally enable our commitment to him or her, reasons *of* love are the reasons a lover takes him- or herself to have with regard to the beloved as the particular person he or she is.

In the second chapter, we will consider the phenomenon of valuing. I will explain why different individuals reasonably value different objects in different ways, and why one and the same object may be valued by some and not by others. The distinction introduced in the first chapter, between reasons *to* love and reasons *of* love, is extended so as to explain the variety of objects and ways of valuing. Individuals stand in different relations to different objects and therefore have different reasons *for* valuing, while a person's particular responses to an object he or she values are properly explained by the reasons *of* valuing the person takes him- or herself to have with regard to it.

Furthermore, I argue that while we may have reason to be practically and emotionally responsive even to objects we do not value, we have *singular* reasons—reasons one has as the particular person one is—only with regard to the objects we value. The final idea introduced in the second chapter is that proper valuing may require devotion, that is, attachment that is not overly concerned with the reasons that support it and that is, at times, blind to reasons against it. Part of our commitment to what we value

is a commitment not to desert it whenever there seems to be sufficient or even conclusive reason to do so.

This somewhat paradoxical aspect of valuing requires some illustration. Consider reasons to be devoted to justice, for instance by struggling to make one's society more just. It might seem that if one's struggle for justice is bound to come to naught, then one lacks reason to be devoted to justice in this way. But devotion to justice might require that we reject the idea that justice is impossible. Martin Luther King was not making a historical claim when he declared that the arc of the moral universe is long, but it bends toward justice. Rather, he expressed his commitment to justice, his faith that justice will prevail despite evidence to the contrary. John Rawls shared this faith, as seems clear from the last paragraph of his introduction to the paperback edition of *Political Liberalism*:

If a reasonably just society that subordinates power to its aims is not possible and people are largely amoral, if not incurably cynical and self-centered, one might ask with Kant whether it is worthwhile for human beings to live on the earth? *We must* start with the assumption that a reasonably just political society is possible, and for it to be possible, human beings must have a moral nature, not of course a perfect such nature, yet one that can understand, act on, and be sufficiently moved by a reasonable political conception of right and justice to support a society guided by its ideals and principles. (Rawls 1993/1996, lx, my italics)

The force of Rawls's "must" is not the force of empirical evidence that justice is possible and that human beings have a moral nature, nor is it a psychological force, compelling us to make such assumptions regardless of whether they are justified. Rather, the locution expresses the force of faith, the strength of a commitment to justice that cannot be undermined by external considerations. This is an important aspect of valuing to which I wish to draw attention, and it also plays a role in the third and last chapter of the dissertation.

The third chapter is twice as long as each of the first ones, partly because it is about the great dark side of love and valuing combined: regret and loss. I argue that we have genuine reasons for regret and discuss three major kinds of regret: regret about past mistakes, regret about significant loss, and regret about one's attachments. I claim that our reasons for regret are importantly dependent on the attachments we have in the present. A change in a person's attachments may give her reason to affirm or celebrate an action that, at the time, she had most reason not to undertake. I also argue that we may regret a significant loss involved in a justified choice without wishing we had chosen otherwise. And, finally, I argue that we may affirm the objects of our attachments while regretting some of the necessary conditions for their existence. In this context, I discuss political predicaments where we seem to be implicated in wrongs committed by institutions we value. In particular, I discuss crimes committed against the Palestinians during the founding of the State of Israel and under Israel's present military rule. I claim that Israelis' devotion to Israel should lead them to fight for its moral worth and legitimacy, to struggle to make it a worthwhile, justifiable object of attachment. But if such revision is impossible, if the wrongs in question are essential to what Israel is, then Israelis should forgo their attachment to Israel, and they should do so with a grave sense of loss.

The overarching theme of the third chapter concerns the relation between one's understanding of oneself as an individual person and one's regrets. The attachments that shape our lives and that in light of which we understand who we are are often the bases of our regrets, whether they are attachments we engage with in the present or attachments we no longer engage with but still entertain. For instance, we may give up a pursuit and

yet continue to value it, and its significance for us may be manifested in the reasons for regret or sense of loss we take ourselves to have with regard to it.

But I also argue that one and the same person may have very different attachments in different moments in his or her life. In this sense, one's self is not defined or constrained by one's present attachments, but reaches beyond them. To be sure, there is a sense in which we are what we love and value, as Brodsky suggests, but there is also a sense in which we are always more than that. The tensions between these two notions of self reflect the tensions between our reasons *for* attachment and our reasons *of* attachment. But the distinction between the two notions of self also makes sense of why one must acknowledge both kinds of reasons in order to be one self.

Chapter I: Loving

Introduction

Like digestion, love is not voluntary. We have very little direct control over who we love and how. Indeed, we often speak of love as a great force that pushes us around. We fall in and out of love, we are blinded and possessed by love, and our love may hold us hostage. Such descriptions bring to mind romantic love, but they apply to love between friends and family, too. Although we can put ourselves in situations that are likely to inspire or stifle our love, we cannot choose to love as we can choose to raise our hand: we cannot love at will.

And yet we are normally more intimately and directly accountable for our love than we are for our digestion. Our love is responsive to our beliefs about our beloveds, about ourselves, and about what we have reason to do, believe, desire, or shun. Sometimes a person's love may seem inappropriate, or misguided, and other times we may find the absence of love objectionable, as in the case of a parent who does not love his child. Even in the grip of love, we can still ask ourselves whether our love makes sense, whether it is right, whether we should give ourselves to it or resist it. It is in this sense that love—but not digestion—is rational: love reflects the agent's judgments about the world around her and may be criticized, endorsed, or disowned.

Nevertheless, the idea that love is rational gives rise to apparent puzzles. First, consider the many children that are in greater need of help than your own child. Is it not objectionable to favor your child over those other children, whose plight is more urgent?

The equal moral worth of all persons seems to be in tension with the disproportionate attention and care we direct at our beloveds. This is the puzzle of *legitimate partiality* (section 1). Second, in loving a person we direct our attention and care to the particular person she actually is, not to some subset of her properties. And yet if love for a person is to be rationally justified, it must be justified on the basis of the beloved's properties. This gives rise to a puzzle about the *particularity of love* (section 2). Third, in loving a person we experience many of our reactions to him or her as called for or required. And yet we may, at the same time, concede that we are not rationally required to love this person. How might love be both required and optional? This is a puzzle about the *necessity of love* (section 3).

The thesis of this chapter is that distinguishing between two kinds of love-related reasons dispels the aforementioned puzzles. On the one side of the distinction in question there are reasons *to* love a person. These are reasons that justify, render appropriate, or make sense of a person's love for another and of actions done from love. On the other side of the distinction there are reasons *of* love for a person. These are reasons that we take ourselves to have insofar as we love a person, e.g., reasons to care about his or her mood, reasons to go hiking together, or reasons to carefully consider his or her opinion.¹ I will argue that, in one way or another, the three puzzles arise from a failure to distinguish and understand the relation between reasons *to* love and reasons *of* love.

¹ We might note a third category, namely, that of reasons *for* love in general, i.e., reasons for forming loving attachments in general. Thus, reasons *for* love would explain why it is reasonable to love anyone at all. After all, some have argued that love is "a losing game." See Amy Winehouse in her album *Back to Black*, Island Records, 2006. http://youtu.be/nMO5Ko_77Hk It is arguable that our reasons *to* love a particular person may be derived from our reasons *for* love in general. I do not think this possibility is in tension with anything I will go on to say.

1. A Puzzle about Partiality

A court judge summons the next case. The defendant enters and stands before her, silently waiting his turn. As the judge raises her eyes from her papers, she recognizes the defendant as her son. The judge recuses herself from the case.

Surely, the judge realizes that she could not remain impartial while having her son as the defendant. Such inability would be sufficient to justify her decision to disqualify herself. But we can say something stronger. It might be argued that even if the judge were able to remain impartial, and treat her son as she would treat any defendant, being impartial toward her son would have been objectionable. So the judge's predicament is not merely one in which it is difficult, or impossible, to do what she ought to do, but one in which she appears to have conflicting responsibilities: being a good mother and being a good judge seem mutually exclusive in this scenario.

Perhaps in this story the conflict between partiality and impartiality may be resolved by the judge's recusal. But unlike the seat of a judge, which one may opt out of, the claims of morality and justice are unconditionally authoritative. We—mothers, fathers, children, siblings, friends, and lovers—cannot relieve ourselves of our responsibilities as moral agents in the way the mother in our example relieves herself of her responsibilities as a judge. Thus, it might be thought, the conflict between morality and partiality cannot be resolved in the way that the judge's conflict is resolved.

Some have concluded from this conflict that, if partiality is ever to be justified, it must be justified by impartial moral considerations (Railton 1984, Baron 1991). Others

have concluded that since partiality may be justified even when it conflicts with morality, we are sometimes justified in violating the demands of morality, at least as those demands are construed by contemporary moral theories (Stocker 1976, Williams 1981a). But many philosophers argue that the problem at hand does not warrant such extreme conclusions: we can make room for legitimate partiality without grounding it in moral justifications and without restricting the decisive authority of morality (Kolodny 2003, Korsgaard 1996, Raz 1989, Scanlon 1998, Scheffler 2001, Scheffler 2010b).²

The apparent conflict between morality and partiality arises because any plausible view of morality accepts the idea of moral equality, namely, “the idea that everyone counts morally, regardless of differences such as their race, their gender, and where they live” (Scanlon 2005, 3). When we are partial toward someone, we seem to violate this basic moral idea. The judge’s son does not count more than any other defendant merely because he is her son.

But moral equality does not, by itself, preclude partiality. For morality certainly does not require that we make no distinctions between individuals. A person in dire need and a person who is sipping a daiquiri on the beach do not warrant the same moral response. While both persons are morally equal, it does not follow from this fact that we should treat them in the same way. Moral principles apply to everyone equally—they are

² The idea that morality and partiality are independently authoritative and generally compatible may leave room for the possibility that there are irresolvable conflicts between them. Susan Wolf has argued this much in Wolf 1992.

universal in their application and general in their content—but they yield different verdicts, depending on the circumstances.³

Once we see that morality itself properly differentiates between persons on the basis of their features and their circumstance, and that this is perfectly compatible with—indeed, necessary for—moral equality, it becomes clear that the problem of legitimate partiality does not follow from the universality of morality, i.e., the fact that moral principles apply to everyone. Universality is compatible with differential treatment because universal principles prescribe actions based on relevant descriptions and have conditions of application that only some individuals satisfy.

If there is a conflict between morality and partiality it is not due to the fact that morality is universal; rather, it is due to the fact that some forms of partiality are morally objectionable. For example, it is arguable that discrimination on the basis of race or gender is objectionable because it is based on features of individuals that are not morally relevant. Similarly, when being partial toward one's friend or sibling is objectionable, this is because the fact that someone is my friend or sibling is not morally relevant to how I should treat him or her in the relevant circumstances. The judge, for example, would be wrong to acquit her son merely because he is her son. But this is a substantive claim about morality rather than a consequence of a fundamental conflict between morality and partiality.

Philosophers have tried to capture reasons of partiality in a way that makes their

³ I take the distinction between universality and generality from Rawls. Rawls suggests that a principle is general if it can be formulated “without the use of what would be intuitively recognized as proper names, or rigged definite descriptions” and a principle is universal if it holds “for everyone in virtue of their being moral persons.” See Rawls 1971/1999, 131-133/113-115.

universality evident.⁴ Just as every person counts morally, and therefore everyone has reason to help anyone in need, so every person may have reason to be mindful of his or her son more than he or she is mindful of a passerby, a reason to help a friend more than a distant acquaintance, etc. Such reasons apply to all persons alike, but they issue different verdicts for different individuals depending on the relation a person bears to other individuals. Thus, while partiality involves special relations between persons, such relations are not special to specific persons but can be had by anyone and refer to anyone, at least in principle. As such, reasons of partiality are universal and compatible with the idea of moral equality.

However, the proposal that the reasons we have with regard to our loved ones are—like moral reasons—universal, may seem implausible. Indeed, the proposal seems to deny the very phenomenon that gave rise to the puzzle of legitimate partiality to begin with. When we love a person we normally respond to reasons that refer to this person as a particular, not to reasons that refer to the person's properties. For example, when my friend Terry tells me of her trouble, I take myself to have reason to offer her my advice, or to listen and console her. But my reason to do so is not that this is how one should respond to a friend in need; rather, it is that *Terry* is in need and this is how I should respond to *her*. If my reasons to advise, listen, and console Terry are universal, as the solution to the puzzle of legitimate partiality suggests, then there seems to be no room for the partiality characteristic of love.

⁴ I have in mind, in particular, the idea of agent-relative reasons, as developed by Thomas Nagel (1970, 47) and, later, by Derek Parfit (1984, 27) and Samuel Scheffler (1994). Agent-relative reasons are normally understood as reasons that contain an essential reference to the agent, as in the claim that everyone has a reason to console *his* or *her* friend. The distinction has also been criticized, for instance by Korsgaard (1993).

To give another, subtler example, if the judge were to recuse herself upon seeing the defendant merely “because it is him” or “because it is Todd,” then her decision would not seem warranted. Indeed, her recusal would seem arbitrary, perhaps even unintelligible. There must be something *about* Todd, or about the circumstances more generally, that makes sense of, or justifies, the judge’s decision to recuse upon seeing Todd in her courtroom. And, indeed, the judge’s love for Todd, and the special reasons she has with regard to him, *can* be explained and made sense of: Todd, the defendant, is the judge’s son. And yet it seems that the judge’s initial surprise and distress upon seeing the defendant in her courtroom is best described as a response to the fact that the defendant is Todd, not as a response to the fact that the defendant is her son.

Thus, if the proposed solution to the puzzle of legitimate partiality is to be plausible, we must show that it is compatible with love’s partiality. Here we first arrive at the distinction that will reemerge throughout this paper: the distinction between reasons *of* love and reasons *to* love. Reasons *of* love are the reasons a lover takes him- or herself to have with regard to his or her beloved. Such are reasons to be mindful of the beloved, to act for the beloved’s sake, to care for the beloved, but also reasons to be offended when the beloved disregards you, to feel implicated in the beloved’s failures, and perhaps even reasons to assign greater epistemic weight to the beloved’s beliefs.⁵ Loving a person involves taking oneself to be responding to reasons *of* love with regard to him or her.

In contrast, reasons *to* love are given by the facts or features that make sense, justify, or require one’s love for another individual. It is in light of a person’s reasons *to*

⁵ For an argument that friendship involves epistemic partiality, see Stroud 2006.

love that her reasons *of* love make sense to us. Examples of possible reasons to love are the beloved's qualities; the kind of history one has with the beloved (Kolodny 2003); the fact that the beloved reciprocates one's love (Brown 1997); or the sheer fact of the beloved's personhood or humanity (Velleman 1999, Setiya 2014).

The puzzle of legitimate partiality gets a grip when we focus on reasons *of* love and neglect reasons *to* love. By themselves, reasons *of* love seem arbitrary, for they refer to the beloved without reference to the features that make him or her an appropriate object of love. But reasons *to* love do just that: they make sense of our love for a particular person and of our reasons with regard to him or her in particular. My reasons *of* love for Todd, for example, make sense in light of the fact that Todd is, say, very kind and wise; or in light of our relationship (Kolodny); because Todd reciprocates the attention I grant him (Brown); or because of Todd's personhood (Velleman). These different answers need not be mutually exclusive. That I have reasons of love with regard to Todd rather than with regard to the table might be explained by Todd's personhood, for example. But were I to explain why I have reasons of love with regard to Todd but not with regard to Terry, I might appeal to the fact that Todd and I have known each other since childhood whereas I only met Terry last week. Different specifications of reasons *to* love assuage different worries about the arbitrariness of reasons *of* love.

Though we are often unsure how to adjudicate the claims of love and morality, there is no reason to think that love is in principle precluded by morality. For if partiality is grounded in genuine reasons (to love), then morality need not be insensitive to the significance of partiality. Indeed, the moral idea that persons as such deserve not to be

treated in certain ways only makes sense when we recognize the legitimacy of persons' partial concerns and attachments. Thus, morality and partiality are not merely compatible; they are complementary.

When the judge recuses herself because the defendant is her son, she does not fail in her role as a judge. On the contrary, she abides by her role as a judge by appropriately disqualifying herself from the case on the basis of the defendant's relevant features (namely, his being her son). After all, as I mentioned above, to save her son she could have resumed her role and ruled in his favor, thereby abusing her authority as a judge. But just as the judge's devotion to the execution of a fair trial is restricted by her devotion to her son, so her devotion to her son is restricted by her devotion to the execution of a fair trial. More generally, moral equality may be compatible with partiality if we can say something general and illuminating about when partiality is permissible or warranted and what are genuine reasons for assigning special significance to specific individuals. To be sure, I have not offered such an account here; I have only claimed that there is no general reason to believe that such an account cannot be provided.

2. A Puzzle about Particularity

In the previous section I argued that the puzzle of legitimate partiality arises when we fail to acknowledge the role of reasons *to* love. Even if loving another partially consists in taking oneself to have reasons *of* love, which refer to the beloved and not to his or her relevant properties, it is our reasons *to* love this individual that make sense of the fact that we have such reasons *of* love for him or her. As long as we have genuine reasons to love a person, favoring them in the relevant ways does not conflict with the

equal moral worth of all persons.

But if the puzzle of legitimate partiality results from overlooking reasons *to* love, our next puzzle, the puzzle of love's particularity, arises from the opposite mistake: neglecting reasons *of* love. The puzzle of particularity begins with the thought that rational love must be grounded in the beloved's properties. The fact that Todd is kind and wise, or that he is my friend, makes sense of my love for him, of the things I do for him and the attitudes I take towards him. But when I help Todd move to a new apartment, for example, I do so out of concern for Todd, not out of concern for kindness, wisdom, or friendship. Insofar as I love Todd, my responses to him are primarily responses to him as the particular person he is. It might therefore seem that if love is rational it must be concerned with the beloved's relevant properties and that this is incompatible with genuine love, which is concerned with the beloved as the particular person he or she is.

In an attempt to clarify this peculiar puzzle, I would like to consider the way in which it emerges from our previous puzzle, the puzzle of legitimate partiality. I begin with an example offered by Charles Fried and famously picked up and criticized by Bernard Williams, of a man who chooses between rescuing his wife and rescuing a stranger from drowning. Fried arrives at the example through a discussion of resource allocation. He considers the question of why we should give priority of resources to actual and present sufferers over absent or future ones. He then offers an analogy: "Surely it would be absurd to insist that if a man could, at no risk or cost to himself, save one of two persons in equal peril, and one of those in peril was, say, his wife, he must treat both equally, perhaps by flipping a coin" (Fried 1970, 227).

Why does Fried think that coin tossing would be “absurd” in such a circumstance? Suppose the husband flips a coin and consequently rescues his wife rather than the stranger. Suppose further that he later justifies his decision to someone by enumerating the morally relevant considerations in favor of flipping a coin. Still we may ask the husband: “What about the fact that she is your wife, didn’t this fact play any role in your decision?” Presumably, if the fact that the woman he saved is his wife played no role in the husband’s decision, we would find it hard to believe that the husband loves his wife at all. It is absurd that a loving husband would act in such circumstances based solely on moral considerations.

Fried proposes a solution:

Where the potential rescuer occupies no office such as that of a captain of a ship, public health official or the like, the occurrence of the accident may itself stand as a sufficient randomizing event to meet the dictates of fairness, so he may prefer his friend, or loved one. Where the rescuer does occupy an official position, the argument that he must overlook personal ties is not unacceptable. (Ibid.)

According to Fried, space can be made for legitimate partiality toward one’s wife by showing that partiality may have a legitimate moral function. In this way, moral demands do not eclipse the husband’s love for his wife.

But to avoid the absurdity in question it is not enough to explain why it is permissible for the husband to rescue his wife; we must also explain why it would be inappropriate—perhaps even impermissible—for the husband to ignore the fact that the woman is his wife. The husband’s attachment to his wife is not merely an emotional or attitudinal state that morality should allow; it is a normative relation to be understood in terms of the husband’s reasons. That is why the husband may wrong his wife by treating her as if she were a stranger.

Thus, even if, in line with Fried's suggestion, the husband considered the fact the accident occurred as a randomizing event that renders rescuing his wife fair, the husband would seem oddly alienated from his wife. "It might have been hoped by some (for instance, by his wife)," Bernard Williams memorably noted, "that [the husband's] motivating thought, fully spelled out, would be the thought that it was his wife, not that it was his wife and that in situations of this kind it is permissible to save one's wife" (Williams 1981a, 18). The idea seems to be that a loving husband whose wife is in mortal danger cannot—indeed, ought not—take up an impartial point of view.

I have suggested that, contrary to Williams's claim, the question of permissibility is pertinent to the husband's choice. Granted, the husband should not ponder the issue and articulate his justification in the moment of action, but the issue is a real one nonetheless. That is, there is a real question about whom to rescue and how rescuing the wife might be justified to the stranger. Furthermore, it is arguable that cases like this one illustrate the authority of morality: morality draws us out of our particular point of view in order that we recognize the particular points of view of others. Not only is it not an objection to impartial morality that it is in tension with the husband's love for his wife; this tension is precisely why morality must be insisted upon.

Still, I think the example and the objection Williams raises reveal more than the inherent tension between love and impartial morality. Williams's point is striking, not because it shows that complete commitment to morality involves inevitable alienation from one's loving relationships, but because it reveals a puzzle about the kind of significance we attribute to the people we love. What is it about the fact that the woman

is the rescuer's wife that makes us shudder when he fails to mention this fact as his reason to save her? Or, put another way, why should the husband be inclined to save her "because she is his wife"?

As I already mentioned, one might not love one's wife, and even if one does love her, it seems odd that one should help her on the basis of one's marital status. A similar thought seems to lie behind Derek Parfit's remark (which Liam Murphy reports): "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 saved her because she was Mary, or Jane, or whatever" (Murphy 2000, 140 n. 36). Admittedly, the very ending of Parfit's remark—"or whatever"—seems to undermine his point, but the idea is that for the husband to focus on his relationship to his wife is to risk being alienated from what ought to be the immediate object of his concern: namely her, Mary.⁶

Nevertheless, it is not clear why the fact that the woman is Mary should provide the husband with a normative reason to rescue her rather than the stranger. Niko Kolodny argues that "[t]he thought that she is Mary simply identifies a particular with itself; it does not ascribe a property to that particular that might make a certain response to it appropriate. After all, the stranger left to drown might point out that he is Fred" (Kolodny 2003, 159). Kolodny goes on to argue that if the fact that the woman is Mary is not a reason to save her, then it cannot be rationally motivating. Kolodny believes, like many others, that rational action is motivated by the agent's grasp of his or her normative reasons. Since rigidly individualized thoughts do not provide normative reasons, Kolodny

⁶ Here I draw on Kolodny's explanation of Parfit's point. See Kolodny 2003, 158.

concludes that the husband could not have been rationally motivated to save the woman “because she is Mary.”

The problem seems to be the following. We have been assuming that the lover’s reasons are like impartial reasons in that they universalize over all agents but pick out specific agents by their properties. Mary must have some property that gives the husband special reason to save her, a property that, in principle, could be instantiated by any agent, not only Mary. Whatever this property might be, it is in virtue of grasping the fact that Mary instantiates it that the husband should see himself as having special reason to save her and therefore be rationally motivated to save her.

Parfit’s remark suggests that even the husband’s recognition of such a property would be one thought too many. The current problem is not that morality is in tension with love; it is that rationality is in tension with love. To rescue your wife because she is your wife is already to be drawn away, or alienated, from one’s particular point of view toward one’s beloved. The husband might protest, for example, that unlike the person who takes the role of the judge and therefore acts in this capacity, he did not save Mary in the role of a husband who saves his wife or the role of a lover who saves his beloved. Mary’s importance to him is not mediated by any general description that he and she meet. Put metaphorically, but adequately, I think, we might say that in rescuing Mary the husband acted as the person he is underneath all his roles and descriptions. And then we should add that he rescued Mary as the person she is underneath all her roles and descriptions. Hence the minimally informative thought: “It’s her!”

To be sure, we need not hold that the husband rescues his wife as the “bare

particular” she is underneath all her properties. A love for a bare particular is highly *impersonal*, for it involves none of the properties that make a person the person she actually is. Rather, what we want to say is that the husband rescues his wife as the particular person she *actually* is, with all of the properties she actually has. What we deny is that in rescuing his wife the husband is concerned with the properties that render rescuing her appropriate.

Whether or not we agree with Parfit’s point in this particular instance, it seems to reflect a powerful intuition, namely, that love involves valuing the beloved as the particular person he or she is rather than as a person who satisfies some relevant description or criterion that warrants love. And the question raised by Kolodny’s resistance to this thought is this: How else are we to make sense of the particularity of love?

*

The question of the particularity of love has been construed in different ways, a fact that, I believe, has led to some confusion. Joseph Raz, for example, thinks of the matter in terms of the beloved’s irreplaceability:

There is (or was) something about the object [i.e., the beloved] which lends it value of a special kind, such that while some feasible replacements may be as good or even better, they will not be quite the same—not quite the same in what makes them good or valuable, and in the precise way that they are or were good or valuable. (Raz 2001, 25-26)

Raz assumes that if the beloved is irreplaceable, then there must be some property that renders the beloved unique in its value for the lover. Of course, the relationships we have with friends, siblings, parents, children, and lovers might be uniquely valuable as a matter

of fact: there might be no one else, or it might be very unlikely that there will be someone else, with whom we can have a relationship that bears the same valuable properties (Raz 2001, 24). But what is needed, Raz believes, is an explanation of how a beloved might be in principle irreplaceable, not merely contingently irreplaceable. So *de facto uniqueness*, as Raz calls it, does not suffice for the irreplaceability characteristic of many loving relationships.

Raz's solution is to appeal to valuable properties that could only be instantiated once: "the first child was the parents' first, and that makes the attachment special, gives it a flavour no other can have for them" (Raz 2001, 27). There is necessarily only one person who gave me my first kiss, one person who is my oldest sister, etc. Raz believes *logical uniqueness*—i.e., the bearing of properties that are by definition satisfied only once by one particular—explains why, in love, we are often attached to a particular person who is in principle irreplaceable. The value of "historical properties," Raz says, "can capture the sense in which what is uniquely valuable is the object—it is the object under the historical description: my first child, i.e., my child qua first child, etc." (Raz 2001, 28 n. 14).

The last comment, about the uniqueness of one's first child, suggests that Raz's explanation misses the mark. For it seems highly implausible that the only thing that renders the parent's child irreplaceable is that she was their first child—as if without this historical property the child might have been replaceable. Reducing the unique value of the child to an historical property is no different from reducing her unique value to some non-historical, repeatable property. For in either case we fail to account for the fact that it

is *this* child that is uniquely valuable, not the properties she instantiates. Raz's account makes it seem as if what we really care about when we see our beloved as irreplaceable are the properties that make him or her so.

Raz rejects this interpretation of his view. He emphasizes that it is the beloved we love, not the features that make him or her valuable. But he says this is not enough to solve the problem of uniqueness, because we love this particular beloved for some reason or other. Raz concludes that since reasons are universal, we must offer a general feature that makes this particular beloved unique (Raz 2001, 28 n. 14).

But from the fact that we love this particular for some reason it does not follow that there must be something about this particular that no other particular has. The same feature that gives us reason to love Jane in particular might give us reason to love Oscar in particular. And the mere fact that both Jane and Oscar have this feature does not entail that one might be replaced by the other. After all, we love the particulars themselves, not their features. So the fact that other individuals who share our beloved's features are available elsewhere need not render our beloved replaceable. Similarly, even if our beloved is logically unique, to use Raz's phrase, it is *she* whom we value uniquely, not the historical property that makes her logically unique.

The confusion arises from the fact that we may have the very same reasons to love various different individuals. But in loving a person we recognize and act on a wide range of reasons that are not our reasons *to* love her: they are our reasons *of* love for her. Unlike reasons *to* love, reasons *of* love essentially refer to a beloved *as a particular*. Raz's mistake, I believe, is that he assigns reasons *to* love the role properly assigned to

reasons *of* love. That is to say, Raz attempts to rationally explain why we love a person uniquely by appealing to our unique reasons *to* love him or her—e.g., that she is my first child, my childhood friend, my first kiss, etc. However, we are normally justified in loving a person uniquely even when we have the very same reason to love another person. Reasons *of* love—not reasons *to* love—explain our love for particulars as such.

Raz is not alone in assigning reasons *to* love a task they are unfit to perform. Since the distinction between reasons *to* love someone and reasons *of* love for him or her has not been made explicit, philosophers have often been led to believe that when we have the same reason to love different persons we cannot rationally love each as an irreplaceable particular.⁷ Before I consider why this distinction should be so elusive, I would like to examine two other examples that manifest philosophers' discomfort with general accounts of the significance of particular persons as such.

Niko Kolodny has developed a thoughtful and elaborate theory of love by examining the vulnerabilities of alternative theories. One such alternative theory is the quality theory. On this theory, our reasons for loving a person are constituted by his or her loveable qualities: his or her beauty, wit, sense of humor, etc. These are traits that warrant our love. One of the objections Kolodny raises against the quality theory is that it fails to account for the “nonsubstitutability” of the beloved:

If Jane's qualities are my reasons for loving her, then they are equally reasons for my

⁷ Roger Lamb is an exception. Lamb calls out and denies the assumption that “if an attitude is universalizable, then its object is in some way a *universal!*” (Lamb 1997, 41-42, italics in original). Similarly, Elizabeth Anderson (1993, 9) and David Velleman (1999, 364) both draw on the Kantian distinction between price and dignity to explain how the capacity for valuing can explain the significance of particulars as such. Unlike things that have only a price, valuing beings have dignity: a valuing being can only be appropriately valued as irreplaceable and incomparable.

loving anyone else with the same qualities. Insofar as my love for Jane is responsive to its reasons, therefore, it ought to accept anyone with the same qualities as a substitute. But an attitude that would accept just as well any Doppelgänger or swamp-Jane that happened along would scarcely count as love. (Kolodny 2003, 140-141)

In contrast to the quality theory, Kolodny's own view, the relationship theory, holds that "my reason for loving Jane ... is my relationship to her: that she is my daughter, or my mother, or my sister, or my friend, or the woman with whom I have made my life" (Kolodny 2003, 146). Furthermore, on this view, Kolodny's love for Jane partly consists in, and is causally sustained by, his recognition that his relationship to Jane renders his love appropriate.

Kolodny believes the relationship theory explains nonsubstitutability: twin-Jane would not warrant his love because she would not be the person with whom he made his life. And though Kolodny agrees that we might imagine a "relationship Doppelgänger," who has the same relational features as my beloved, he believes that in this case it is indeed appropriate to love the Doppelgänger:

If my wife and I decide to have a second child, for instance, then we bring into this world a relationship Doppelgänger to our first child. The relationship theory implies that we have just as much reason to love the second child as the first. But this is the right implication. (Kolodny 2003, 147)

Though it is certainly plausible that Kolodny and his wife have just as much reason to love their second child as they have to love their first, it is utterly implausible that their second child may be substituted for their first. But if Kolodny's complaint against the quality theory is valid, then he is committed to the claim that his children are substitutable. Recall that, according to Kolodny, the quality theory fails to account for nonsubstitutability because it implies that Kolodny has the same reason to love Jane as he does to love twin-Jane: both Jane and twin-Jane are witty, kind, funny, etc. But

Kolodny's own view implies that he has the same reason to love his second child as he does to love his first: both are his children. Kolodny is therefore similarly exposed to the objection of substitutability. But as we have seen from our discussion of Raz, the right response on behalf of both the relationship theory and the quality theory is that we may have the same reason *to* love different individuals as the particular persons they are and therefore have reasons *of* love not to substitute one for the other. From the fact that we have the same reason to love two different individuals it does not follow that one may replace the other.⁸

Finally, consider Harry Frankfurt's view, which denies that there are reasons *to* love a particular person. Frankfurt takes the particularity of love as his primary starting point and observes, like Raz and Kolodny, that love involves a principled resistance to substitution: "Substituting some other object for the beloved is not an acceptable and perhaps not even an intelligible option" (Frankfurt 1999, 166). Furthermore, Frankfurt holds that "a person cannot coherently accept a substitute for his beloved, even if he is certain that he would find himself loving the substitute just as much as he loves the beloved that it replaces" (Frankfurt 1999, 169).

But unlike Kolodny and Raz, Frankfurt takes for granted that any account of love in terms of value or reasons would fail to do justice to the particularity of love. He argues

⁸ Perhaps a more charitable interpretation of Kolodny's complaint against the quality theory is that the quality theory yields implausible conclusions about what facts constitute reasons to love, not that it is committed to the substitutability of the beloved. This interpretation would enable Kolodny to say, against the quality theorist, that it is implausible that I have the same reason to love twin-Jane, whom I just met, as I have to love Jane, with whom I made a life. Moreover, it is plausible that Kolodny has the same reason to love his second child as he does his first. But all this is besides the present point, which is that, like Raz, Kolodny slides back and forth between implicitly accepting and denying that general features of particular persons may account for our love for them as the particular persons they are.

that “the reason it makes no sense to consider replacing what we love with a substitute is not that loving something entails supposing that it is one of a kind” (Frankfurt 1999, 169). Since Frankfurt assumes that reasons could only account for the beloved’s irreplaceability if the beloved were “one of a kind,” he concludes that any reasons- or value-based view of love would fail to capture the sense in which the beloved is irreplaceable. It is not anything about the beloved that renders him or her unique: “The focus of a person’s love is not those general and hence repeatable characteristics that make his beloved describable. Rather, it is the specific particularity that makes his beloved nameable” (Frankfurt 1999, 170).

These observations lead Frankfurt to the conclusion that love consists in a complex structure of desires, or, as he puts it, “a complex volitional structure that bears both upon how a person is disposed to act and upon how he is disposed to manage the motivations and interests by which he is moved” (Frankfurt 1999, 165). Frankfurt seems to think that desires, unlike values and reasons, can take particulars as such as their object. Desires therefore seem to him fit to explain the particularity of love.

But, as we have seen, from the fact that the object of love is the particular beloved as such it does not follow that love cannot be grounded in “those general and hence repeatable characteristics that make [the] beloved describable.” To say that my love for Yaara is warranted by her sensitivity, and wisdom, and passion for life and beauty, is not to say that if another person with these qualities came along she would render Yaara replaceable; nor is it to say that in loving Yaara I only pay attention to these qualities; nor is it to say that I love Yaara as a way of loving someone who has these qualities. To

borrow a distinction offered by Kolodny, Yaara's qualities may provide the *ground* of my love for her but it is Yaara who is the *focus* of my love (Kolodny 2003, 154). There may be universal reasons for loving a particular person as such. Contrary to what Frankfurt seems to believe, and notwithstanding the many merits of his view, the particularity of love is compatible with the idea that love is responsive to reasons.

*

One possible source of confusion about the particularity of love may be an implicit assumption about the content of the lover's reasons—namely, that reasons to respond to the beloved must include the beloved's general features, not merely a reference to the particular person. More specifically, the assumption is that for the beloved's properties to be part of the rational explanation of the lover's reasons, the beloved's properties must be part of the lover's reasons. This assumption leads to the implausible conclusion that the lover's reason—for example, to care and act for the sake of her beloved—includes the properties that make love rationally appropriate. But this is not a plausible account of the content of reasons *of* love. What we want to say, instead, is that, as in the case of the husband who rescues his wife, the lover's reasons are normally not the properties that make love appropriate, but the beloved him- or herself, as well as the relevant circumstances.

Fortunately, it has been argued that the assumption that yields the implausible consequence is false. Jonathan Dancy, for example, distinguishes between facts that favor an action and facts that enable other facts to favor an action (Dancy 2004, ch. 3); T. M. Scanlon distinguishes between a consideration that counts in favor of an action and the

conditions under which the consideration counts in favor of an action (Scanlon 2014, ch. 2); and Mark Schroeder distinguishes between a reason for action and the background conditions in virtue of which a proposition counts as a reason for action (Schroeder 2007, ch. 2).

The idea is that there may be normative explanations of why we have the reasons we have that are not part of the reasons themselves. The fact that I pursued a degree in philosophy was a condition for, or enabler of, my reason for writing a chapter in my dissertation about rationality. If I did not pursue a degree in philosophy, I would not have had reason to write a chapter about rationality. But the fact that I did pursue a degree in philosophy was not itself a reason to write a chapter about rationality. Whether a person has a reason to perform a given action normally depends on certain background conditions, but these background conditions are not part of the agent's reason to perform the action.

It therefore seems possible to maintain that what explains the lover's reasons are the beloved's relevant properties and at the same time deny that these properties are part of the lover's reason for performing the action. The husband's reason for rescuing Mary is simply that Mary is in mortal danger, but Mary's other properties—e.g., her being his wife for the past 25 years—partly explain why the husband has this reason. The distinction between *reasons of love* and *reasons to love* may be understood along the lines of the distinction between reasons and background conditions, or favorers and enablers, to use Dancy's parlance. Admittedly, as we fall in love, or have a "crush," we respond to the features of the person that warrant our love, to our reasons *to* love him or

her. But once in love, the very same features recede to the background and resurface only in moments of reflection or conflict.

When the two kinds of love-related reasons are run together, love might seem preoccupied with its own appropriateness. Niko Kolodny, for example, believes that love

partly consists in the belief that some relationship renders it appropriate, and the emotions and motivations of love are causally sustained by this belief... Special concern for a person is not love at all when there is no belief that a relationship renders it appropriate. (Kolodny 2003, 146)

It seems to me that Kolodny's proposal does not do justice to the particularity of love. Even if Kolodny is right in holding that love is rendered appropriate by the kind of relationship one has with one's beloved—i.e., even if he's right that relationships give us reasons to love a particular person—it is implausible that special concern can only count as love if the agent believes that her love is made appropriate by her relationship with the beloved. It is not uncommon to love without knowing why, or how to make sense of it. For special concern to count as love it is sufficient that the agent takes herself to have reasons *of* love with regard to the person toward whom the concern is directed: reasons to spend time with the person, to mind his or her thoughts and feelings, etc. If we come to believe that we lack reasons *to* love, or that we have reasons *not to* love, then we might reassess our reasons *of* love, or face a genuine conflict between reasons that pull us in opposite directions. But the question of the appropriateness of our love for a particular person might not arise, and we may therefore not have a clear view about the matter, even as we are deeply in love.

Kolodny's claim that what sustains love is primarily a belief in its appropriateness also seems inaccurate. Just as we might be troubled by the impression that our love is

rationality criticizable, or objectionable, or that it does not make sense, so our love might compel us to reassess our views about whom it is rational to love. This is not to deny that our beliefs about appropriateness are part of what sustains our love, but to point out that our love is part of what sustains our beliefs about appropriateness: the two are in reflective equilibrium.

Furthermore, love is often appropriately oblivious to the question of its own appropriateness, as illustrated by the case of the husband who rescues his wife without thinking. And even when a person recognizes that she has very strong reasons, or overall reason, to love someone, it would be odd if her love depended on this belief. Indeed, contrary to Kolodny's contention, I suspect that a love that relies on its appropriateness to survive is no love at all.

3. A Puzzle about Necessity

We have thus far dealt with two puzzles about the rationality of love. With regard to the puzzle of legitimate partiality, we saw that if there are genuine reasons to love, then the actions and motivations of love need not be objectionably arbitrary or conflict with morality. In considering the puzzle of love's particularity, we saw that our reasons to love a particular person are not the reasons we primarily respond to in loving them. Our reasons, e.g., to mind the beloved's feelings, to consult in him or her, to act for the beloved's sake, and to spend time together, are reasons *of* love. Reasons *of* love refer to the beloved as a particular, not merely as a person we have reasons *to* love. The particularity of love is thus salvaged by the distinction between two kinds of love-related reasons.

But a significant challenge remains. Consider the following question: does it follow from the fact that I have reason to love a person that I am required to love her? A positive answer would seem more plausible in the case of familial love than in the case of friendship or romantic love. The fact that a child is my daughter may be said to give me conclusive reason to love her. But in the case of friendship or romantic love, love seems essentially optional: having reasons to love someone as a friend or lover does not imply that it would be inappropriate or rationally criticizable not to. Love may often be rationally permissible but not required.

That love and the actions done from love may be optional in this way would not have seemed puzzling if it weren't for the fact that in loving a person we see our emotions, as well as many of our actions, as called for or required. The beloved's good, his or her thoughts, feelings, and concerns, carry significant weight in our deliberation and play an important motivating role in our lives. And though love certainly involves more than positive feelings, and may often inspire annoyance, anger, and even rage, love always assigns a special authority to the beloved without which neither our happiness nor our anguish make sense.⁹ The problem is that the fact that love is peremptory seems compatible with the idea that often love is not rationally required. Oddly, the emotions of love, and the actions and reactions to which they give rise, appear to be both required and

⁹ David Velleman complains that recent accounts of love in analytic philosophy “express a sentimental fantasy” in which “love necessarily entails a desire to ‘care and share,’ or to ‘benefit and be with’. But,” says Velleman, “it is easy enough to love someone whom one cannot stand to be with” (Velleman 1999, 353). In agreement with Velleman's general observation, I would like to add that while our feelings and attitudes toward people we love can range from rage and anger to warmth and affection, the one attitude we cannot take toward our beloveds is indifference. This is why, when faced with a choice between being hated and being ignored, more often than not people choose the former. Anonymity might seem worse than notoriety.

optional. This is the puzzle about *the necessity of love*.

It seems plausible that the two kinds of love-related reasons we have been discussing correspond to the optional and necessary aspects of love: the optionality of love is accounted for by reasons *to* love, while the necessity of love is explained by reasons *of* love. Since we can have reasons *to* love someone without having reasons *of* love for him or her, love may be optional. But when we have reasons *of* love for someone, love is preemptory.

I think this is the right thing to say here, but it is not enough in order to disarm the puzzle about love's necessity. For now we must explain how there could be reasons *to* love a person but no reasons *of* love with regard to him or her. The worry is that if reasons *of* love *follow* from reasons *to* love, then we have reasons *of* love with regard to anyone we have reasons *to* love. But the proposed account of love's necessity turns on the possibility that we can have reasons *to* love a person without having reasons *of* love with regard to him or her. So in order to address the puzzle of love's necessity, it is not enough to draw the distinction between the two kinds of love-related reasons—we must also say something about the relation between them.

We may begin to do this by noting that love is not alone in its predicament of “optional necessity.” Promising, too, is optional in one sense and necessary in another: making a promise is optional, but keeping it is not.¹⁰ Just as we have reasons *to* love and reasons *of* love, we may be said to have reasons *to* promise and reasons *of* promise. Like reasons *to* love, reasons *to* promise are general and often optional, and like reasons *of*

¹⁰ I thank Aleksy Tarasenko-Struc for this point.

love, reasons *of* promise refer to particulars as such and are peremptory. It might therefore be suggested either that we promise to love or that love is a kind of promise. This would explain why we have optional reasons *to* love, but peremptory reasons *of* love.

But while it is certainly true that loving a person normally involves making promises to him or her, and that such promises might further and deepen our love, I do not believe that the necessity of love can be appropriately construed as the binding force of a promise. Promises are only binding when they are made voluntarily, but love is often involuntary. And, unlike the case of promising, even when one acknowledges the involuntariness of one's love, this does not diminish the apparent force of one's reasons *of* love.

Christine Korsgaard has urged me to consider the fact that the traditional Christian marriage ceremony has people promising to love, that in Victorian novels, young men ask women to love them, and that people say things like "I give my love to you freely."¹¹ These observations suggest that there is a case to be made for the voluntariness of love after all, and that loving might be closer to promising than I have thus granted.

But while Korsgaard is surely right that promises, and the idea of voluntary love more generally, have always played a crucial role in the practices of love, I believe this role must be understood in light of the fact that love itself is *involuntary*. For unlike a failure to keep a promise, which gives rise to a claim against the promisor, when love

¹¹ Personal correspondence.

dissipates, the lover cannot be accused of violating a voluntary commitment. To be sure, we can promise to do all in our power to keep our love alive; we can promise to be there for someone come rain or come shine; and we can promise to be honest about our love and not misrepresent it; but love cannot be reduced to a promise to do certain things, nor can we genuinely promise to love.

Furthermore, it is striking that in the three examples offered by Korsgaard, the voluntariness of love is *declared*. Why should voluntariness be declared? Consider the fact that in order to establish that we can voluntarily raise our hand, there is no need to seek declarations of voluntary hand raising. I suspect that declarations of this sort count against the voluntariness of love rather than for it. Just as vehement denials of guilt may inadvertently reveal it, emphatic declarations of voluntary love may often indicate love's *involuntariness*.

“I give my love freely,” “I promise to love you forever,” “I give you my heart”—perhaps, rather than expressions of the lover's choice to love, these declarations are meant to enable the lover to guarantee the authenticity of his or her love. Of course, the authenticity of any declaration of this sort is itself open to doubt, but its goal may be the removal of doubt, all the same. It is not an exaggeration to say that the history of love is the history of doubts about love's authenticity.¹² It would therefore make sense that over

¹² In the story of the Binding of Isaac, the testing of Abraham may be seen as an expression of God's doubt with regard to Abraham's love. Part of what is striking about the story is Isaac's *lack* of doubt in his father's love. Isaac famously says to his father: “Behold the fire and the wood; but where is the lamb for a burnt-offering?” (*Genesis*, 22:7). The very existence of God's doubt in Abraham renders Isaac's trust in Abraham misguided. Moshe Halbertal argues that the Binding of Isaac is an expression of God's “rich-spouse predicament”: “He can never be sure that he is not loved for his money, or to put it more bluntly, he will always be in doubt as to whether he is loved altogether” (Halbertal 2012, 22).

the course of human history, people have devised various ways to “confirm” and “secure” their own love as well as the love of their beloveds. The futility of such “confirmations” does not reduce their urgency, nor does it undo the longing that drives us to assimilate loving to promising.

A second answer to the puzzle of love’s necessity avoids claims about the voluntariness of love and insists that reasons *of* love do follow from reasons *to* love. It may be argued that while a person has some reason to love almost anyone, she has overall decisive reason to love only a few. Some facts or considerations may count in favor of loving a person and yet be outweighed or defeated by other facts or consideration. In contrast, when we have overall decisive reason to love, all the relevant considerations have already been taken into account and the final verdict requires our love. On this proposal, then, we have *reasons of love* only with regard to the few individuals whom we have *overall decisive reason to love*. What makes love seem optional, on this account, is that we normally have reasons to love without those reasons amounting to *overall decisive* reason to love, and therefore without having reasons *of* love.

My main worry about this proposal is that it is implausible that we are rationally *required* to love the persons who are our friends or romantic partners. No matter how wonderful someone might be, no matter how loving and caring, even if you enjoy this person’s company, find him or her interesting, attractive, and trustworthy, and you have a long, wonderful history with him or her, you may not take yourself to have reasons of love with regard to this person. And I do not think it is plausible to say that you would

thereby be open to rational criticism for failing to recognize your reasons *of* love. For you might simply not have reasons of love with regard to this highly loveable person. It may be rationally permissible not to love him or her: sometimes love is not forthcoming and that is all there is to it.

A third proposal might seem more plausible. It may be argued that while we may not be rationally required to love a person, once we actually love someone, and have reasons *to* love this person, then we have further reasons *of* love with regard to him or her. Love would then consist of a set of attitudes, motivations, and emotions that take a particular person as their object and, in the appropriate conditions, give rise to a wide range of reasons with regard to the beloved. Like the promise-based account, this proposal also appeals to a fact about the agent that triggers reasons of love; but it does not rely on love's voluntariness. The idea, then, is that the emotions, attitudes, and actions of love are optional before the fact of love and non-optional once love is in place.

But such an explanation of reasons *of* love puts the cart before the horse, so to speak. For no behavior, psychological state, sensation, or disposition can count as love unless it reflects the agent's judgment that she (already) has reasons *of* love with regard to the particular beloved. We might put the point by saying that in loving a person we normally see ourselves as answerable to him or her in particular. That is to say, loving someone involves taking that person to have a kind of entitlement to make certain demands on us and to hold ourselves accountable for not meeting certain expectations. We are answerable in different ways to our lovers, friends, parents, children, and siblings. There are many (many) things that we do not share with our children or that even our

very good friends do not know about us, but we are still answerable to our children and friends in the ways that are relevant to our love for them. No matter how we feel about a person, if we do not see him or her as significant to us in this peculiar way—if we do not see ourselves as somehow bound up in the person’s particular point of view—we do not love the person. But if we count as loving a person in virtue of seeing him or her as authoritative—i.e., in virtue of taking ourselves to have reasons of love with respect to the person—then our love cannot *explain* the beloved’s authority and our reasons *of* love.¹³

The foregoing observation suggests that our reasons *to* love another, and our inclination toward loving him or her, are not sufficient to give rise to reasons *of* love. It seems to me that in order to find the missing ingredient that explains reasons of love, we should look, not at the lover, but at the beloved. Whether a person has the authority characteristic of a beloved depends not only on our emotions and attitudes toward him or her and our reasons to have those emotions and attitudes, but also on the emotions and attitudes of our would-be beloved. Our reasons *to* love may provide a rational basis for wishing that the other person would see us as his or her beloved. But only if the other person in fact reciprocates our plea and accepts the authority we wish to grant him or her, do we have reasons *of* love with regard to this person. And while a person can only have reasons *of* love with regard to another if she also has reasons *to* love him or her, her

¹³ There’s a wrinkle here. We can see our beloveds as authoritative independently of whether they are in fact authoritative. So it may be argued that love is constituted by the fact that a person sees another as authoritative in the relevant way and that seeing someone as authoritative *makes* that person authoritative. Love may therefore explain reasons *of* love. But it would follow from such an explanation that there could be no justification for taking oneself to have reasons *of* love other than that one already takes oneself to have reasons *of* love. And this seems odd, to say the least.

reasons *of* love crucially depend on the beloved's uptake.¹⁴ As such, reasons *of* love entail reasons *to* love but are not entailed by them, for they are essentially inter-subjective. This is the thesis about *the inter-subjectivity of reasons of love*: they are reasons that depend not only on reasons to love, but on reciprocal willingness to love.

It follows from this thesis that even when we have reasons *to* love that rationally permit but do not require our love, we may have reasons *of* love due to our beloved's stance towards us. Such a predicament explains why it makes sense to experience oneself as failing to appropriately love a person even when one is not rationally required to love him or her. For example, you might feel that you are drifting away from a friend, or a lover, but that you should continue to love him or her. Perhaps one day, out of the blue, you find yourself unmoved by your beloved's sadness. Or your beloved tells you how they feel about you, that you are their dearest friend, or that they never felt so close to someone, and you realize that you cannot truthfully say the same about them. And when you realize this you may very well acknowledge that there are others whom it would make sense for you to love—i.e., that there are other loves you may rationally pursue instead—and that by ceasing to love this person you would not be violating any promise you made to him or her. You may realize that you are not generally required to love this person. And yet you may be determined to resist the undoing of your love due to your commitment, due to your beloved's authority, due to your reasons *of* love. The beloved may seem to have a claim to your love, and you may seem to remain answerable to him or her, even if your attitudes and emotions pull you away and toward someone else.

¹⁴ The solution I propose here is akin to the promise-based solution insofar as promises, too, require uptake on behalf of the promisee. But, in contrast to the promise-based solution, on the solution I am proposing neither the offering of love nor its acceptance need to be voluntary.

The opposite might happen, too. You may wish not to love someone anymore, and vigorously pursue other relationships or friendships that you have reason to pursue. And yet later—maybe years later—you may come to realize that all along you have been failing to acknowledge that the person you left behind is important to you. Perhaps you have ceased to love that person and haven't given them much thought over the years. Perhaps you made the right decision overall in moving on. But now, as you are reminded of your old beloved, you realize that in navigating your love away from him or her you acted against your reasons *of* love.

In Marcel Proust's *In Search of Lost Time* we find a poignant example of the way a person may come to recognize his or her reasons of love long after the feelings of love have dwindled or grown numb. The narrator has scarcely mourned his grandmother's death, but more than a year later the pain of grief suddenly overwhelms him:

I clung to this pain, cruel as it was, with all my strength, for I realized that it was the effect of the memory I had of my grandmother, the proof that this memory was indeed present within me. I felt that I did not really remember her except through pain, and I longed for the nails that riveted her to my consciousness to be driven yet deeper. I did not try to mitigate my suffering, to embellish it, to pretend that my grandmother was only somewhere else and momentarily invisible... Never did I do this, for I was determined not merely to suffer, but to respect the original form of my suffering as it had suddenly come upon me unawares, and I wanted to continue to feel it, *following its own laws*, whenever that contradiction of survival and annihilation, so strangely intertwined within me, returned. (Proust 1938/1999, 215, my italics)

Suffering has the authority of law in Proust's depiction of bereavement. The narrator's suffering derives its authority from being an expression of his grandmother—of her presence and existence. Grief descends upon us when we feel, all at once, the authority of the lost-beloved and its absence; the “contradiction of survival and annihilation.” In the depths of grief, the possibility that our pain will subside, that we will be able to do without our beloved, seems like a second killing, both of our beloved and of the person

we were when we were with him or her. Living on without this person may seem as a denial of their particular authority, of their indispensability to us, and therefore as an admission that, all along, we were able to do without him or her. And even if we have most reason to move on, to accept our beloved's absence, we also have reason never to let go of the vivid memories and the inevitable pain that accompanies them. These memories and pain function as a fragile vessel through which our long gone beloved remains present. That there are often reasons—albeit, desperate reasons—to hold on to those we loved, and that these reasons have a crucial role in a lover's grief, suggests that the attitudes, motivations, and emotions of love reflect judgments about reasons *of* love.¹⁵

Furthermore, reasons of love do not only address particular persons as such, but particular persons as they actually are. Consider the thought that things could have worked out very differently; that some of the most important people in our lives are people we could have just as easily never met. And yet, as things actually turned out, these people are indispensable to us: the thought of losing them is unbearable and the thought of never having met them is startling. The same is true of one's children, who may have not been born at all, or may have been born at a different time and therefore as different children. We may recognize that had we had our child a few months later, the hypothetical, younger child would have been important to us in the same way that our actual, older child is important to us. But the thought of this alternative history makes us cringe because for us as we actually are, no other child would do. From our actual perspective, not having this child would have been a grave loss, even if it were no loss at

¹⁵ In the third chapter of this dissertation, I develop an account of rational regret along these lines. I argue that regret and a sense of loss are warranted by our attachments to persons and things we failed in some way and yet continue to love and value.

all for the person we would have been.

This thought is eloquently expressed in Robert Musil's novella "The Perfecting of a Love," in which a woman, Claudine, grapples with the extent and boundaries of her love for her husband. At one point, Claudine reflects: "Somewhere among all these people there is someone—not the right one, someone else—but still, one could have adjusted oneself even to him, and then one would never have known anything of the person that one is today" (Musil 1986, 215). The hypothetical beloved Claudine contemplates is not "the right one" for her, but he could have been the right one—or could have seemed to be the right one—for the person she would have been. Claudine would have been a different person had she had reasons *of* love with respect to someone else.

Thus, if I love Yaara then I see myself as answerable in relevant ways to her as she actually is. The fact that I recognize that it would have been rationally permissible for me not to love her, or to love someone else instead, does not settle the question: Am I answerable to Yaara? And I cannot settle this question by asking: Do I love Yaara? For my answer to the latter question is at least partly determined by my answer to the former: I love a person only if I take myself to have reasons *of* love with regard to him or her. If I appropriately see myself as answerable to Yaara, this is because Yaara accepts my wish to love her, and reciprocates it.

*

We have considered different ways of construing the relation between reasons *to*

love and reasons *of* love so as to explain the sense in which love is both optional and necessary. One solution maintains that loving is akin to promising, or to promising one's love. On this account, love is optional before its promise and non-optional after. Another solution holds that when a person has overall decisive reason *to* love another, she has reasons *of* love with regard to him or her. Love would then seem optional when we have *some* reasons to love that do not amount to *overall decisive* reason to love. The third proposal distinguishes between optional reasons *to* love before the fact of love and non-optional reasons *of* love after the fact of love. Once we love someone appropriately, this proposal maintains, we have reasons *of* love with regard to him or her.

My own solution to the puzzle of love's necessity distinguishes between the optionality of reasons *to* love and the non-optionality of reasons *of* love. I have suggested that the fact that we have reasons to love someone and that we wish to love this person cannot settle the question of whether we have reasons *of* love with regard to him or her. Our reasons *of* love do not only depend on our reasons *to* love and our wish to love, but also depend on the corresponding attitude of our beloved. I have reasons of love with regard to Yaara partly in virtue of *Yaara's acceptance* of the authority I wish to grant her, partly in virtue of *my wish* to grant her such authority, and partly in virtue of *my reasons to* grant her such authority. From the point of view of my reasons *of* love with regard to her, my love is preemptory, but from the point of view of my reasons *to* love Yaara, my love is merely permissible. There are general facts about the conditions in which it is rationally appropriate or even required to love another. But whatever those conditions are, the normativity of love does not follow from them, for it concerns the authority of actual, particular individuals considered as such. Digestion concerns nothing of the sort.

Chapter II: Valuing

Introduction

The people we love fill most of our waking hours as well as our dreams, but so do the things we value. A healthy, meaningful life includes activities, pursuits, practices, and ideals to which one devotes oneself, as well as appreciation and enjoyment of things engaged with, pursued, or produced by others. We come across landscapes that inspire our admiration; we encounter artistic and scientific accomplishments that invoke our esteem; we are astounded by nature's creations and we exalt manifestations of moral virtue, beauty, and wisdom. We also find ourselves offended and beleaguered by vulgarity and baseness, while the destruction of things we hold dear and the failure of our pursuits may leave us inconsolable. The current chapter finds its point of departure in the thought that while the persons we love bestow meaning and a sense of worth on the lives we lead, our lives revolve to no lesser degree around the things we value.

At first glance, the rationality of valuing may seem less puzzling than the rationality of love. As we have seen, a primary problem about the rationality of love is that love appears to defy reason in its insistence on particularity. We love particular persons as such, not merely as persons we have reasons to love. Valuing, however, does not seem to have this problematic aspect, for valuing presents itself as an exercise of reason. Valuing an object—be it a goal, an activity, an ideal, an accomplishment, or an artifact—seems to involve responsiveness to the object's apparent value. While we can reasonably love someone without having any sense of why we should love him or her in

particular, it seems that to value an object we must judge the object valuable in a way that makes sense of our stance toward it.

What remains puzzling about valuing is that there are many valuable things we recognize as such but do not value, or do not value to the degree we value other valuable things, while the things we value most need not be the ones we deem most valuable. Furthermore, one and the same thing might be properly valued in different ways by different individuals. In other words, there seems to be a break between our objective assessment of an object and our subjective attachment to it.

Thus, just as we do not expect others to love our beloveds or to love them as we do, we do not expect others to value what we value or to value it as we do. Israeli society, the Hebrew language, Tel-Aviv, Jerusalem, Israeli literature and culture—these are all things I value but do not expect most of the readers of this dissertation to value in any way or degree similar to the way and degree I value them. Philosophy, on the other hand, is something all of us—you, the readers of this dissertation, and I, its author—value in similar, if somewhat different ways. Nevertheless, most people do not value philosophy as we do, and this does not render them unreasonable or mistaken. (Thus, it goes without saying that this dissertation is not intended as a bestseller.)

To explain the rationality of valuing we must make sense of the disproportionate significance we attribute to the things we value and of the corresponding fact that many valuable things leave us cold. That is to say, we must account for the personal significance of what we value in a way that reconciles it with the requirements of rationality. This is what this chapter aims to accomplish.

In the first section I characterize the phenomenon of valuing and argue that an individual's relation to the objects he or she values is part of the rational explanation of his or her valuing. In the second section I consider the claim that we normally and reasonably value only a fraction of the things we believe valuable. I argue that the gap between valuing and believing valuable is not best understood in terms of our emotional vulnerability to an object. Rather, I argue in the third section, the gap should be understood as the difference between taking oneself to have only universal reasons, reasons that everyone has, and taking oneself to have also singular reasons, reasons that are particular to oneself.

In the fourth and final section, I discuss a remaining, endemic puzzle. I argue that often proper valuing requires a degree of obliviousness to one's reasons for and against valuing. We may have reason for commitment or devotion to a person, ideal, vocation, tradition, or goal, and therefore have reason to occasionally disregard the rational basis of our devotion. Alternatively, when we recognize that we no longer have reason to hold on to an object we value we may experience this very realization as a failure of our devotion to it. These tensions and conflicts are an ineliminable part of human life, and should be acknowledged as such. Thus, the thesis of this chapter as a whole is that different individuals may rationally value one and the same thing in different ways, and one and the same valuable object might be rationally valued by some and not by others, and yet, often, in order to value properly we must disregard, deny, or defy the rational basis of our devotion.

1. What Is Valuing?

In order to get a better sense of the phenomenon, let me start with some examples of things we value. Relationships are a major example in the philosophical literature. The most common example of a kind of relationship that may be reasonably valued is friendship, but other kinds of relationships are mentioned as well, such as familial relationships or romantic relationships (Kolodny 2003, 2010; Wolf 2010; Wallace 2013; Raz 2001; Scanlon 1998). Indeed, it has been argued that we value particular individuals as a way of valuing our relationships with them (Kolodny 2003), but some insist that we may value particular individuals as such and independently of the relationships we share with them (Setiya 2014; Keller 2013). Then there are examples of valued “projects.” Such projects may involve either production of some thing or state of affairs or engagement of some sort in certain activities. Studying philosophy and practicing the cello are examples of projects in which one engages (Wolf 2010, 4). On the other hand, succeeding in one’s career, possessing the entire set of nineteenth-century French stamps, or qualifying as an International Master in chess are examples of projects that involve production (Raz 1999b, 63). So are climbing Kilimanjaro (Nagel 1986, 167) and writing a good book on Kant’s ethics (Korsgaard 1993). Still other examples of valuing involve appreciation, support, or preservation of different sorts: valuing a friend’s sense of humor, valuing paintings, valuing literary genres, valuing historical artifacts, valuing one’s privacy (Scheffler 2010a, 18-22), valuing the US Constitution (Scanlon 1998, 95), and valuing All Souls College (Cohen 2013). Finally, we may value certain ideals, practices, activities and traditions both in the sense that we engage them and in the sense that we appreciate them and wish to preserve them (Scheffler 2010b). Some of the

authors of the examples above did not intend to explore the notion of valuing as such, but all the examples fall under what is generally meant by this label and illustrate the phenomenon I wish to explain.

With these examples in mind, we may make the following observations. First, valuing an object involves taking a certain point of view on it. From this point of view, one adopts and experiences certain attitudes and emotions with regard to the object one values, and one takes oneself to have certain reasons for action with regard to it. Whether I value climbing Kilimanjaro or the US Constitution, my valuing involves a range of reactions and attitudes, as well as ascription of deliberative significance to considerations relevant to the object I value. Thus, valuing is not an individual attitude, action, or choice, but an orientation that includes various attitudes, actions, and choices.

Accordingly, when we speak of reason *for* valuing we do not speak of reason to take a certain action with regard to an object, but of reason to orient oneself in a certain way with regard to it. By contrast, our particular actions with regard to the object we value, as well as our particular attitudes and emotions with regard to it, are explained by the reasons *of* valuing we take ourselves to have with regard to it.

One's reason *for* valuing may be part of the rational explanation for one's particular reasons *of* valuing. For example, the fact that the painting is beautiful might give someone reason *for* valuing it, and this reason may partly explain why the destruction of the painting gives the person a reason *of* valuing to be sad and feel a sense of loss. The distinction between reasons *for* valuing and reasons *of* valuing brings to the

fore the difference between valuing as a general orientation and the particular attitudes, actions, and emotions that are constitutive of valuing.

My second observation concerns what may be called the difficulty of valuing. Unlike choosing and doing, valuing is not voluntary. This, of course, is not to say that there cannot be reasons to value something, for even our non-voluntary attitudes may reflect or commit us to judgments about reasons (Scanlon 1998, 19-21). But it is important that valuing is not the kind of thing we can directly decide upon. Valuing the goal of climbing Kilimanjaro is not the same as deciding to climb Kilimanjaro. One may climb, or at least try to climb, Kilimanjaro even if one does not value doing so. Similarly, and perhaps more plausibly, one can play the cello without valuing playing the cello. Admittedly, it may be argued that in order to play the cello well one must value doing so, but then playing the cello well is not something one can directly decide upon either.

This last example suggests that the difficulty of valuing goes rather deep. Even when we are non-voluntarily inclined to value an object we believe we have reason to value, such inclination is normally not sufficient to enable us to value it. In order to value something we must become responsive and attentive to it in certain non-voluntary ways. Wanting to take on such an orientation toward a valuable object is still a long way from actually taking it on.

This second observation, about the difficulty of valuing, leads to the third: although we cannot voluntarily value, and we are normally unable to value even when we sincerely and rationally want to, we can *learn* to value. Practice, or, as Aristotle would put it, habituation, shapes the way we see things, it molds our non-voluntary responses

and adjusts the range of our voluntary responses—it changes our character. Thus, in enabling us to do things we were not able to do and respond to an object in ways we would not have responded to it otherwise, habituation enables us to take points of view that were not available to us before. We overcome the difficulty of valuing by learning to value.

For a long time I wished to value philosophy by engaging in it. I was enticed by books I read, experiences I had, and, later, by introductory courses in moral philosophy, ancient philosophy, epistemology, and metaphysics. But despite my sense that there is something fascinating and truly valuable about philosophy, I did not feel that I could properly engage in this activity, that I could adopt the way of thinking that is distinctive to philosophy. I appreciated various arguments that I had encountered and the insight they offered, but could not understand how philosophers come up with such arguments and insights; I gained a better understanding of myself and of the people around me by studying philosophy, but did not know how to fruitfully engage the ideas I encountered. I struggled to get a handle on the kind of deliberation that philosophy requires. It took a long while and a whole lot of practice before I sensed that I might—perhaps, possibly, conceivably—value philosophy properly, that I might be able to *do* philosophy.

Valuing something properly normally requires practice, not only in order to succeed in doing and feeling what one believes one has reason to do and feel with regard to it, but also in order to recognize those reasons. In other words, learning to value involves more than learning *how* to value; it involves learning *what* valuing consists in. Another semi-autobiographical example might prove helpful here. When I first learned to

play the upright bass, my teacher gave me some basic rules about how to position the instrument, how to hold the neck of the bass and spread my fingers, how to move the bow on the strings, etc. I remember that after a few bass lessons, I came across a video recording of a famous bass player, Charles Mingus. I was amazed to see that he did not follow any of the basic rules given to me by my teacher. I reasoned that since Mingus is a well-known bass player, whose excellence is unquestionable, my teacher's rules must be mistaken and my teacher cannot be a very good one. It did not occur to me that in order to learn how to play a musical instrument one must follow rules that are not themselves essential to playing the instrument but function as scaffoldings that habituate the student and enable him or her to develop the sensitivities and skills that *are* essential to playing it. Once one *knows* how to play the bass, it no longer matters how one positions the instrument or how one grips the neck or the bow. It was precisely because Charles Mingus was a great player that he did not need to follow the valid educational rules my teacher gave me.

Furthermore, habituation is required even when the object one values is not something with which one engages, but rather something one appreciates or admires. In listening to classical music or hip-hop, reading literature, or viewing contemporary art, one learns to value through experience. I came to see the various virtues that make Mozart a great composer, or Notorious B.I.G. a groundbreaking hip-hop artist, by listening to and learning about classical music and hip-hop. To value something properly, one must develop a taste for it, a capacity to judge appropriately with regard to it.

Even the most common forms of valuing require learning. As Aristotle observes: “All men enjoy in some way or other good food and wines and sexual intercourse, but not all men do so as they ought” (Burnyeat 1980, 77). And M. F. Burnyeat argues that for Aristotle, habituation is crucial not only in order to *do* what is noble, good, or virtuous, but also in order to *know* what is noble, good, or virtuous:

I may be told, and may believe, that such and such actions are just and noble, but I have not really learned for myself (taken to heart, made second nature to me) that they have this intrinsic value until I have learned to value (love) them for it, with the consequence that I take pleasure in doing them. To understand and appreciate the value that makes them enjoyable in themselves I must learn for myself to enjoy them, and that does take time and practice—in short, habituation. (Burnyeat 1980, 78)

Thus far I have suggested that to value something we must inhabit a certain point of view with regard to it, a point of view from which we may see and feel facts about it that call for the responses and actions that constitute valuing it. Furthermore, I suggested that learning to value involves shaping our attitudes, motivations, dispositions, and emotions so as to recognize and properly respond to reasons *of* valuing with regard to the object we value.

The final observation I would like to make concerns the difference between *valuing from afar* and *valuing up close and personal*. Valuing an art form, for instance, might mean little more than appreciating and enjoying it occasionally, thereby *valuing it from afar*. But it might also mean being a connoisseur of the art form, practicing it, or devoting oneself to it in some other way, thereby *valuing it up close and personal*. Normally, we understand from the context which sense of valuing is intended. When you say you cannot stand opera and I say that I value it, I probably mean that I appreciate opera, not that I am an avid operagoer. However, when I say to my friend that I value our relationship, when I say that I value my hometown, or when I say that opera is one of the

things I value most in life, I do not merely suggest that I appreciate and enjoy these things occasionally, but that they play a significant role in my life, that they are especially important to me and that they make my life meaningful.

One and the same valuable thing may be valued in different ways by different individuals: some value it from afar while others value it up close and personal. To make sense of the different ways of valuing it is not enough to appeal to the value of an object, we must also attend to the relation between the particular agent and the object in question. Thus, the way I value the town of Ness-Ziona cannot be understood independently from the fact that this is the town where I spent my childhood. Similarly, when I say to the friend who proclaimed her aversion to opera that I value it, it may be clear to her that I am not invested in opera. My friend understands that I am not a composer, an opera singer, nor an opera connoisseur, and that in saying that I value opera I merely mean that I take a generally appreciative stance toward it.

The relational aspect of valuing is twofold. Our reasons *for* valuing something in one way or another are given by facts about our relation to it. Such reasons explain our emotional, deliberative, and attitudinal orientation to the object in question, as well as whether it makes sense for us to value it from afar or up close and personal. But our specific reactions to the object we value are explained by the reasons *of* valuing we take ourselves to have with regard to it. The fact that Ness-Ziona is my hometown explains my reasons *for* valuing it as I do, but the fact that suburban neighborhoods have replaced the orchards of my childhood explains my reason *of* valuing to lament the change Ness-Ziona has undergone. Thus, different ways of valuing one and the same thing are

explained by different reasons *for* valuing it as well as by different reasons *of* valuing with regard to it.

We now have an apparent explanation of why one and the same valuable object can be reasonably valued in very different ways by different individuals. Insofar as the value of the object is understood in non-relational terms, it is only part of the rational explanation of its significance to a particular individual. Another important part of the explanation is the relation between the individual and the object. Once relevant facts about this relation are brought into view, we can see that some individuals have reason for valuing that other individuals lack, and that different kinds of valuing may be called for and constituted by different reasons *of* valuing with regard to the valuable object.

2. Valuing and Believing Valuable

In his insightful and influential essay on valuing, Samuel Scheffler notes: “it is not only possible but commonplace to believe that something is valuable without valuing it oneself” (Scheffler 2010a, 21). As examples of activities he recognizes as valuable but does not value, Scheffler gives folk dancing, bird-watching, and studying Bulgarian history. Scheffler seems reasonable in valuing only a fraction of the activities he takes to be valuable, and so do we.

The distinction between reasons *for* valuing and reasons *of* valuing may explain differences between ways of valuing, but it does not explain why it may be rational not to value objects we judge valuable. Suppose one does not value opera, neither from afar nor up close and personal. One simply does not enjoy it at all. It seems that such a person

might still believe that opera is valuable for its own sake, unlike, for instance, counting blades of grass or rolling a large stone up and down a hill. What is the rational basis of such indifference toward that which one deems valuable?

Scheffler believes the gap between valuing and believing valuable is at least partly explained by the fact that valuing involves emotional vulnerability: “most people have a capacity for the recognition of value that far outstrips their capacity for emotional vulnerability” (Scheffler, 2010a, 27). He goes on to argue that valuing any X involves at least the following elements:

1. A belief that X is good or valuable or worthy,
2. A susceptibility to experience a range of context-dependent emotions regarding X,
3. A disposition to experience these emotions as being merited or appropriate,
4. A disposition to treat certain kinds of X-related considerations as reasons for action in relevant deliberative contexts. (Scheffler 2010a, 29)

First, I should note that it is quite common to experience context-dependent emotions with regard to the object one values without experiencing them as merited or appropriate. I might believe I should not care about a certain place, activity, or relationship, and yet value it very much and exhibit a range of context-dependent emotions with regard to it. To be sure, if I value the object in question despite my judgment that I ought not value it, then I exhibit a rational failure, but I value it nonetheless. I agree with Scheffler, however, that the emotions constitutive of valuing normally express or commit us to certain normative judgments. Such emotions are essentially intentional, that is, they are *about* the object we value and therefore reflect a certain understanding of its significance. One can value even when one believes one has most reason not to value, but valuing involves taking oneself to have reasons *of* valuing.

But whether or not one agrees with Scheffler's characterization of the emotions constitutive of valuing, the list above offers, or seems to offer, an explanation of the difference between valuing and believing valuable. Valuing something, according to Scheffler, involves much more than a mere belief that it is valuable. It is therefore not surprising that we value only a fraction of what we believe valuable. We may have reason to be emotionally vulnerable only to a few of the things whose value we recognize.

To see the force of Scheffler's point, consider an apparent analogy between value and truth. One might assume that valuing is to value as belief is to truth. Insofar as a person is rational, when he or she forms the belief that *p* is true, he or she believes *p*. Similarly, we might initially think that insofar as a person is rational, when he or she forms the belief that that there is conclusive reason to believe *X* is valuable, he or she values *X*. Scheffler's point is that this is not the case. To be sure, if we come to believe that there is conclusive reason to believe that *X* is valuable then we should, as a matter of structural rationality, believe that *X* is valuable, but believing so does not commit us to valuing *X*. This is because valuing goes well beyond believing valuable, according to Scheffler, and involves emotional vulnerability.

However, Scheffler also argues that believing something is valuable commits us to taking ourselves to have certain reasons with regard to it. Even if, like Scheffler, we do not value folk dancing, bird-watching, and the study of Bulgarian history, the fact that these are valuable activities implies that "we should not, for instance, cast aspersions on those who engage in the activities or disrupt what they are doing without good cause" (Scheffler 2010a, 34). Those who do value these activities, Scheffler says, have

additional reasons that nonvaluers, like us, lack. But believing that something is valuable commits us to recognizing certain minimal reasons with regard to it.

Scheffler's claim about the minimal reasons of nonvaluers raises a problem for his account of the gap between valuing and believing valuable. The problem is that the minimal reasons Scheffler invokes seem to include reasons to be emotionally vulnerable. If the minimal reasons of nonvaluers include reasons for emotional vulnerability, then nonvaluers are valuers after all: believing valuable commits one to valuing. But are we really committed to emotional vulnerability with regard to anything valuable?

It seems to me that we are. When we witness the humiliation of a stranger, the oppression of a culture that is not our own, or the destruction of an important work of art that we do not enjoy, we might feel apprehension, anger, indignation, or sadness, and such emotions and attitudes may be called for by our reasons with regard to the object in question as well as by our reasons with regard to the persons involved. More generally, our disposition to avoid harming or disrupting valuable objects, relationships, persons, projects, and activities, is properly coupled with a tendency to experience certain emotions when such disruptions and harms occur. A valuable object need not play a significant role in one's life in order for one to have reason to be emotionally responsive to it in certain contexts. I need not enjoy opera to regret that the only opera house in town is closing down. Scheffler's account of valuing, when combined with his plausible claim about the reasons we have with regard to anything valuable, implies that there should be no gap between valuing and believing valuable: we should value all that we believe valuable.

We arrive at the following *reductio ad absurdum*:

1. Believing something is valuable does not commit us to valuing it. (Scheffler's gap between valuing and believing valuable)
2. Valuing X involves at least believing that X is valuable, being emotionally vulnerable with regard to X, believing that one's emotions with regard to X are merited, and having a disposition to treat various X-related considerations as reasons for action in relevant deliberative contexts. (Scheffler's account of valuing)
3. If something is valuable then we should have certain minimal reactions to it. (Scheffler's claim about the reasons nonvaluers have)
4. So, believing something is valuable commits us to certain minimal reactions to it. (From 3)
5. The minimal reactions we should have with regard to what we believe valuable have all of the elements mentioned in 2.
6. So, having such minimal reactions is a form of valuing. (From 2, 5)
7. So, believing something is valuable commits us to valuing it. (From 4, 6)
8. So, believing something is valuable both does and does not commit us to valuing it. (From 1, 7)

The above argument suggests that we should either reject 1, that is, the claim that we may reasonably believe that something is valuable without valuing it, or reject 2, which is Scheffler's account of valuing. If, as Scheffler holds, valuing involves deliberative and emotional vulnerability with regard to an object we believe valuable, then the claim that we should be deliberatively and emotionally vulnerable with regard to anything valuable implies that we should value anything we believe valuable. Rather than rationally explain the gap between valuing and believing valuable, Scheffler's account leads to the conclusion that there should be no such gap. However, if there is a justifiable gap between believing valuable and valuing, then we should reject Scheffler's account of

valuing as insufficient. There must be something else to valuing that distinguishes it from the orientation we should take toward anything valuable.

The reader might resist this disjunctive conclusion by rejecting 5, the claim that the minimal reactions we should have toward what we believe valuable include emotional reactions. You might wonder how we can reasonably be expected to be emotionally vulnerable to everything that we believe valuable? And you may resort to Scheffler's claim that most people have a capacity for the recognition of value that far outstrips their capacity for emotional vulnerability. However, I believe this claim overestimates our capacity for the recognition of value and underestimates our capacity for emotional vulnerability.

Consider, first, our capacity for emotional vulnerability. While we cannot be emotionally vulnerable in the same way to all that we believe valuable, we may indeed be emotionally vulnerable in restricted ways to various valuable things that do not play a significant role in our lives. Of course, we cannot be emotionally responsive to all that we believe valuable at once. But emotional vulnerability to what we believe valuable only implies that there are certain, specific contexts in which we take ourselves to have reason to be emotionally responsive to what we believe valuable. Complete strangers may warrant our compassion and empathy in certain circumstances, and we may have reason to be saddened by the extinction of a language that no one speaks any longer.

As for our capacity to recognize value, note that recognizing the value of something is closely connected to being emotionally vulnerable to it. When we hear of a disaster that befell people in a far away place, our failure to be emotionally responsive to

the disaster casts doubt on our recognition of the event's significance and of its impact on the lives of the individuals involved. Alternatively, the mere belief that the event is horrible may be seen as a first crack in the dam of one's indifference to it, and, as such, as the *beginning* of one's recognition of its significance. Similarly, I cannot genuinely recognize the value of Bach's Toccata in D Minor BWV 913 if in listening to it I remain utterly unmoved. Even if I can deliver a detailed explanation of Bach's accomplishments in this work, I fail to perceive its value if I am completely indifferent to the music. Again, the mere belief that the composition is valuable might be a first step toward recognizing its value, but it is not more than a first step.

Perhaps, then, we should distinguish between *recognizing the value* of an object and *recognizing that it is valuable*. Suppose that my friend who is a music expert tells me that Bach's Toccata in D Minor BWV 913 is a masterpiece. I then come to believe that it is valuable without ever having heard it and without knowing anything else about it. Indeed, I might have no idea *why* the toccata is valuable, but only believe that it is. My belief might very well be justified, given my friend's expertise. I then recognize that the toccata is valuable without recognizing the value of the toccata.

This might be Scheffler's predicament with regard to folk dancing, bird-watching, and the study of Bulgarian history. And this might also be the reader's predicament with regard to the Hebrew language. You do not recognize the value of the Hebrew language—recognizing its value would involve being responsive to it in ways that you are not—but you may still recognize *that* the Hebrew language is valuable.

The distinction between recognizing that an object is valuable and recognizing its value lines up with the distinction between believing valuable and valuing. However, there is the following difference between the two distinctions. Recognizing that an object is valuable, unlike believing it is valuable, is factive. One cannot recognize an object that is not there, but one may believe a proposition that is in fact false. The same is true of recognizing the value of an object and valuing it. Recognizing the value of an object implies that the object has this value, but one can value an object in a way that fails to reflect its value. The difference is important when we consider the phenomenon from a third-person perspective. The claim that someone values an object does not commit us to the claim that he or she recognizes its value. But the difference disappears from point of view of the agent, for insofar as the agent is (structurally) rational, she takes her belief that an object is valuable to be a recognition that it is valuable, and she takes her valuing of the object to be a recognition of its value.

Admittedly, merely recognizing *that* something is valuable, without recognizing its value, does not commit us to valuing it. But, as Scheffler says, in certain situations one should be deliberately and practically attentive to valuable objects that one does not value oneself. I have suggested that in such situations one should also be *emotionally* attentive to that which one believes valuable, and it is in this sense that one should be emotionally vulnerable to an object even if one does not recognize its value but only *that* it is valuable. Thus, if valuing an object involves believing it is valuable and being emotionally and deliberately responsive to it, then believing that something is valuable commits one to valuing it and recognizing that something is valuable commits one to recognizing its value.

But I think we should instead reject Scheffler's account of valuing. For it seems to me that we want to hold that *recognizing* or *believing that something is valuable* does not commit us to *recognizing its value* or *valuing it*. You may recognize that Hebrew is valuable without being committed to learning the language or engaging with it in any way. We must then hold that, contrary to Scheffler's claim, emotional vulnerability does not explain the gap between valuing and believing valuable, for both valuing and believing valuable call for emotional vulnerability. We should therefore resolve the disjunctive conclusion of the argument laid out above in favor of rejecting Scheffler's account of valuing as insufficient. Now we have the task of amending Scheffler's account of valuing so as to make sense of the fact that we may reasonably believe something is valuable without valuing it, and that we may reasonably recognize that something is valuable without recognizing its value.

An alternative account of the distinction between believing valuable and valuing holds that to value something is to be attached to it. Since we lack reason to be attached to everything we believe valuable, and there are many valuable things we have reason *not* to be attached to, we have reason *not* to value everything we believe valuable. Furthermore, attachment seems to properly separate what we value from what we merely believe valuable. It is arguable that even things we value from afar are things we have a certain attachment to, but valuable things we do not value at all are also things to which we are not attached in any way. This is the difference between the person who values opera from afar and the person who does not enjoy or value opera at all but recognizes that it is valuable. The nonvaluer of opera might still regret the closing of the only opera house in town, but he is not attached to opera and therefore does not value it. It might

therefore be proposed that attachment is the missing ingredient in Scheffler's account of valuing.

However, it is not clear what attachment consists in and therefore what invoking it adds to our account of valuing. Valuing and attachment are so neatly aligned that the two phenomena might be one and the same. If they are, then invoking attachment to account for valuing amounts to giving our problem a different name. Now, instead of an account of valuing, we are in need of an account of attachment. If we want to appeal to attachment in order to offer an informative account of the distinction between valuing and believing valuable, we must explain what is involved in being attached to an object.

Another possible explanation of valuing that aims to make sense of the distinction between valuing and believing valuable, holds that believing valuable is a very weak form of valuing. According to this line of thought, we may describe the meaning-conferring role that valuable objects play in the lives of individuals in order to distinguish between stronger and weaker forms of valuing. We may then locate *valuing up close and personal*, *valuing from afar*, and *merely believing valuable*, on a descending scale from the most intense, demanding, and meaning-conferring form of valuing to the least intense, demanding, and meaning-conferring one.

But there is reason to think that the difference we want to elucidate is not a mere difference in degree. Believing valuable need not commit us to valuing, not even in a minimal way. Still, certain strong emotional reactions toward valuable objects we do not value are warranted when we, for example, see a valuable object being destroyed. Such was the case with regard to the ancient artifacts that were recently destroyed in Iraq by

ISIS. Many of us did not value these ancient artifacts—indeed, most of us were not even aware of their existence before their destruction was filmed and broadcasted worldwide. Nor did we come to value them after the fact of their destruction. But since we believe that these artifacts were valuable, we are appalled by their destruction.

Many valuable objects are, so to speak, none of our business: they do not concern us, and we normally do not care about them in our everyday lives. That is why we are inclined to say that there is a *gap*, rather than a *slope*, between valuing and believing valuable. With regard to many persons and valuable things, we take an impersonal stance of respectful disinterest, which does involve certain obligations and occasional emotional responsiveness but does not involve valuing.

Consider another observation that supports the idea that there is a difference in kind between valuing and believing valuable. An attitude of respectful disinterest is often seen as the appropriate attitude to take toward strangers one comes across in everyday life. The barista who makes your coffee, the bus driver who drives you to work, or the passerby who asks you for directions, may not be individuals you care about or value, but you still have certain obligations to them, you should be courteous and respectful toward them, and you should not be indifferent to their call for help. Now, it is a striking fact about the Israeli public sphere that such a stance of respectful disinterest toward strangers is normally absent from it. An Israeli barista might ask you about your political views while making your coffee, an Israeli bus driver might tell you about his relationship with his wife while giving you change for your fare, and an Israeli passerby might give you advice on how to raise your children while asking for directions. It can be suffocating or

heart warming, but in Israel, for better or worse, individuals who cross paths with you usually care about you, even if they've never met you before and will never see you again. In other words, in Israel even interactions with strangers tend to be *personal*. Insofar as Israelis should adopt an attitude of respectful disinterest toward strangers, they should not merely care less about strangers, or value them less, but adopt a different stance altogether, one that is essentially *impersonal*. In the next section I explain the difference between valuing and believing valuable in a way that accounts for the distinctively personal stance we take toward the persons and things we value.

3. What Makes Valuing Personal?

Thus far I have noted a distinction between reasons *for* valuing and reasons *of* valuing and suggested that it enables us to evaluate, justify, and make sense of different ways of valuing one and the same thing. I then argued that valuing something involves more than taking oneself to have reasons for action and emotional vulnerability with regard to it, and that there is a difference in kind, not merely a difference in degree, between believing valuable and valuing, or between recognizing that something is valuable and recognizing its value. Valuing someone or something implies taking a personal stance toward him, her, or it, whereas believing that something is valuable is fully compatible with a disengaged, impersonal outlook. Next, I wish to consider what accounts for the personal aspect of valuing. We find a promising line of thought in a recent discussion of aesthetic judgment.

Richard Moran discusses the different concepts of beauty we find in Kant and Proust (Moran 2012). Kant believed that a judgment that something is beautiful demands

universal agreement. According to Kant, if I judge that a flower is beautiful, I am committed to holding that all should make the same judgment. This, Moran argues, raises difficulties for Kant, who holds that aesthetic judgment is free in the sense that it is determined neither by our concept of the object nor by our interest, desire, or need (Moran 2012, 315). What, then, is the source of one's authority over others' judgments of beauty? Why should one's aesthetic judgment be determined by another's aesthetic judgment?

Moran argues that in Proust we find a different view of aesthetic judgment that preserves what is attractive about the Kantian view and dispenses with the demand for universal agreement. Both Kant and Proust believe that aesthetic judgment is objective in the sense that it does not merely reflect the agent's attitudes, inclinations, or likings, but is primarily expressive of requirements and demands that the agent might satisfy or fail to satisfy, that is, requirements or demands that are independent of the agent. In this regard, the beautiful contrasts with the merely agreeable. As Moran explains:

[F]or both writers, the idea of something whose status as a value does not depend on my current desires or interests brings to the experience of that value a sense of my being measured by it (rather than my estimating it according to my own needs) and a normative direction of fit from oneself to the beautiful object rather than the reverse. (Moran 2012, 322)

Both Kant and Proust wish to capture the impression that beauty is recognized, not merely felt or experienced, and that one's response to the beautiful is experienced as required or necessitated. The beautiful object *calls for* our acknowledgment.

But while Kant believes that such necessity can only be made sense of as a universal requirement, a requirement that all will acknowledge the beauty in question, Moran's Proust holds that the requirement in question might be singular, that is, a

requirement addressed to a particular agent as such. Thus, on the Proustian view Moran articulates, the sense of necessity characteristic of aesthetic judgments lacks the demand for universal agreement (Moran 2012, 317). The requirement is independent of the agent's interest or desire, and is in this sense objective, but it is particular to the agent, and is in this sense subjective.

Moran compares this notion of singular requirements, as we might call them, to the requirements of love and valuing:

The necessity experienced and lived out in one's relation to a person or a vocation does not translate into a desire for or even the possibility of universal agreement. Obligations and their necessities can come singly and individualize or even isolate the person who finds himself subject to one. (Moran 2012, 318)

Moran is surely right that obligations and necessities with regard to what we value and love come "singly and individualize," and that we should not expect others to face the same necessities as we do with regard to the persons and objects we care for. Nevertheless, there remains a universal aspect to these necessities, for the claim that one faces such obligations and commitments with respect to the objects of one's love and valuing is a claim with universal authority. For example, taking myself to have reason to work all night on this dissertation commits me to the claim that I should work all night on this dissertation, and this claim is a claim that others may agree with or contest regardless of whether *they* should work all night on this dissertation. Claims about a person's obligations to an object are universal even if the obligations they are about are obligations of a particular person as such.

Perhaps, then, we may clarify Moran's point by distinguishing between recognizing that an object is beautiful and recognizing its beauty. We can say that Kant is

right to hold that a true judgment that an object is beautiful is true for everyone and in this sense universal, and yet we may side with Moran and Proust in holding that recognizing and responding to the object's beauty need not involve an expectation that everyone will do the same. Nor does it involve the expectation that everyone will believe that the object is beautiful. There are many true propositions that are not relevant to one's life, or that one never considers; it is perfectly reasonable not to believe all that is true. Our recognition of an object's beauty may be particular to us even if the claim that the object is beautiful is a claim that is true independently of the relation between the object and any particular agent.

Following Moran's lead, we may appeal to the distinction between *universal* and *singular* requirements and reasons in order to illuminate the difference between recognizing that an object is valuable and recognizing its value, or between believing valuable and valuing. When one merely believes that an object is valuable, one is committed to the belief that everyone, including oneself, has reason not to disrupt, harm, or destroy the object. But when one values an object, one sees various facts about the object as reasons for oneself in particular to respond to the object in various ways. In other words, what sets valuing apart from believing valuable is that valuing an object involves taking oneself to have obligations and commitments to it that are particular to oneself and are not shared by everyone who recognizes *that* the object is valuable. The gap that Scheffler explains in terms of emotional vulnerability is more properly described as a gap between taking the fact that the object is valuable, or worthy of valuing, as a reason for anyone to respond to it in certain ways, and taking various facts about the object as reasons for oneself in particular.

Let me clarify the distinction between universal and singular reasons. The distinction is not the same as the distinction between agent-neutral and agent-relative reasons. Agent-relative reasons are normally understood as reasons that contain an essential reference to the agent, as in the claim that everyone has reason to console *his* or *her* friend in a moment of grief.¹⁶ Agent-neutral reasons, on the other hand, lack an essential reference to the agent, as in the claim that everyone has reason to help a person in need. Both agent-neutral and agent-relative reasons are universal in the sense that they are preceded by a universal quantifier: “everyone.” The two kinds of reasons differ in whether they make their conditions of application explicit: an agent-relative reason includes a description of the agent, while an agent-neutral reason does not include such description.¹⁷

In contrast, the distinction between universal and singular reasons is a distinction between reasons that include a universal quantifier, and as such may be either neutral or relative, and reasons that include an essential reference to a particular agent as such. An agent’s reason not to harm a valuable object is a universal (neutral) reason that all agents have, or it is a universal (relative) reason had by all agents that are able to harm the object. But my reason to write this dissertation is a reason only I have, a singular reason. To be sure, my reasons for valuing philosophy as I do are universal in the sense that all agents who stand in a certain relation to philosophy have such reasons. Furthermore,

¹⁶ See my note 4 in chapter 1.

¹⁷ This way of drawing the distinction is akin to Mark Schroeder’s (2007, 18): “the agent-neutral relation is to be analyzed as a universal quantification into the agent-place of the agent-relational relation.” Thus, “for *R* to be a [neutral, O. N.] reason to do *A* is for *R* to be an agent-relational reason for all of [us] to do *A*” (Schroeder 2007, 18). In other words, agent-relative reasons make the description of the agents explicit, while agent-neutral reasons keep the description implicit. On this account, the most neutral reasons are reasons relative to all agents.

these universal reasons for valuing philosophy are part of the rational explanation of my singular reasons to write this dissertation. Still, my reason to write this dissertation is a reason I have as the particular person I am, it is a reason for me in particular. Thus, even if one's reasons *for* valuing something are universal in the sense that anyone in such position has reason to value it, one's reasons *of* valuing include reasons that address the valuer in particular as such and are therefore singular reasons.

Finally, the distinction between universal and singular reasons helps us explain why valuing plays a crucial role in conferring meaning on individual lives. While all persons are equal from the point of view of universal reasons, the reasons special to a particular agent—singular reasons—distinguish him or her from all others. To put the point somewhat metaphorically, we may say that singular reasons give content to a person's bare particularity. Thus, the gap between valuing and believing valuable is essential to our view of ourselves as particular individuals, for it reflects the fact that each one of us is uniquely related and committed to different valuable things.

The above is not meant as a conclusive argument for a specific account of valuing. Rather, I take myself to have made a case for the claim that any plausible account of valuing would hold that valuing involves taking oneself to have singular reasons with regard to the object one values. The gap between valuing and believing valuable is thus explained by the fact that an individual normally has singular reasons only with regard to very few of the things he or she believes valuable.

4. A Lingering Puzzle About Devotion

Different individuals value different valuable objects while one and the same valuable object is normally valued in different ways by different individuals. My goal thus far has been to explain the rational bases of these two facts. In the first section I argued that our varying reasons *for* valuing and reasons *of* valuing rationally explain our different ways of valuing. In the second and third sections I argued that valuing involves taking oneself to have singular reasons with regard to the object one values, and since an individual has singular reasons only with regard to a few valuable objects, an individual may rationally value only a fraction of what he or she believes valuable. In the final section of this chapter I argue that valuing may require obliviousness to one's reasons for valuing and therefore warrant disregard for the rational conditions of one's devotion. This aspect of valuing exposes a tension between the notion of valuing and the idea that we may have reasons for and against valuing.

Consider a vivid example offered by Joseph Raz (1999c) and later discussed by Niko Kolodny (2011).

Let us imagine a ballet dancer enjoying a reasonably successful career with a small provincial ballet company, the only one within hundreds of miles. Then dwindling audiences threaten the future of the company. If it is forced to disband our dancer will have to abandon ballet, to change career and look for something else to do. It does not surprise us that he regards the prospect as a great personal disaster. I think that it would be agreed that it is reasonable for him to try to prevent the collapse... Let us assume that the company has to close. Our dancer looks for other possibilities and starts a new career as a theatre director with the local theatre company where he remains until his retirement. He quickly comes to like his new work, enjoying a success comparable to his success in his first career. (Raz 1999c, 316)

The dancer is devoted to his dancing career, which means that even had he been aware that he will come to value a new career as a theater director, he may still have regarded the threat to his dancing career as a disaster, and reasonably so. We may say that the

dancer has reason to be devoted to dance despite apparent reasons to let go and move on to his next career as a theater director. What we want to understand is what facts give the dancer reason for devotion? Why would it not be more reasonable to give up dancing and embrace his new career as a theater director?

Part of Raz's point in telling the dancer's story is that the dancer's well-being cannot be his reason to try to save his company. The dancer is not concerned with his well-being, for he might believe that his well-being will not suffer from the shift from dance to theater and yet reasonably resist the closing of the ballet company. Although he will become a successful theater director who values his work at least as much as he presently values dancing, the dancer has reason to work hard and sacrifice much in order to maintain his dancing career and stave off the change in his attachments.

It is worth noting that well-being is a bad answer to our question in two distinct ways. First, it is unlikely that the dancer is devoted to dance *in order* to enhance his well-being. This would make the dancer's devotion implausibly instrumental. As Stephen Darwall points out, many philosophical accounts of well-being wrongly assume that "the agent's good is a highest-level, or most final, rational end that structures all of an agent's first-order rational pursuits" (Darwall 2002, 82). Rather, a person's own well-being is one of various rational ends the person might have. We normally value our aims, projects, and loved ones for their own sake, not for the sake of our own good (Scanlon 1998, ch. 3; Scanlon 2010, 94-95; Raz 1999c). As the example of the dancer shows—and many cases of reasonable self-sacrifice demonstrate even more acutely—our own good might come in conflict with the good of what we value (Darwall 2002, ch. 2; Overvold 1980).

But the proposal that one's well-being gives one reason to be devoted to the object one values fails even if it is construed differently. We might invoke well-being not as a final end to which all our other ends are instrumental, but as the rational basis of adopting our ends, as the *ground of* or *reason for* our pursuits rather than their *focus* or *object*.¹⁸ Various facts about the things we value give us reason to value them for their own sake. I admire this flower because it is beautiful, not in order to admire beauty, nor in order to admire beautiful things. The beauty of the flower is the ground of or reason for my admiration, whose focus or object is the flower itself. It is part of what makes the flower valuable, but it is not what I value. Similarly, it may be thought, the fact that the dancer's career enhances his well-being is the ground of or reason for his devotion to his dancing career. But even on this construal, according to which one's well-being is not one's sole rational end but rather the ground of all one's rational ends, the proposal fails. It fails because the dancer seems reasonable in his devotion to his dancing career independently of whether such devotion enhances his well-being. The enhancement of his well-being cannot, therefore, be the reason for his devotion to dance.

Richard Moran offers another case of devotion that he finds in Proust's *Swann's Way*. As Moran explains, Swann suggests to Marcel that

if he moved away from France and formed other attachments he would not be tormented by Gilberte or other unattainable objects of his desire. But the prospect of other attachments, along with the forgetting of his current ones and thus the relief from the suffering they are now causing him are not seen by him as any form of compensation for or solution to his current sufferings but only as adding to them. (Moran 2012, 327)

¹⁸ For the distinction between the two ways in which something may be good in virtue of something distinct from it, see Korsgaard (1983; 1999; 2008), Darwall (2002), and Kolodny (2003, 154).

Like Raz's dancer, Marcel recognizes that he may form new attachments and leave behind his current, unsuccessful ones, particularly his love for Gilberte. But Marcel's devotion to Gilberte is not undermined by this realization, just as the dancer's devotion to his dancing career is not undermined by his realization that he will form an attachment to a new career as a theater director. And even if we grant, as we should, that the dancer and Marcel are reasonably devoted to the objects of their attachments, and that devotion of this sort is essentially non-instrumental, we are still left with a question about the rational basis of such devotion.

Moran expresses sympathy for Harry Frankfurt's answer to this question: "the thought is not that love itself is demanded or required but that loving something or someone *creates* necessities for the person, marking out the limits of what can coherently be willed, creating obligations of responsiveness and attention" (Moran 2012, 318, n. 20). On this, attitude-based view of devotion, the fact that one values a particular object gives rise to reasons of valuing and devotion with regard to it. And though we do not value something for the sake of valuing it, the fact that we value it gives us reason *of* valuing to act for its sake. The resulting answer is seductively natural: our devotion to an object we value is rationally explained by the fact that we value it.

But Moran also argues that the beautiful object, like the objects of love and valuing, differs from the merely agreeable object. Objects we value give rise to "a demand on the subject that is independent of one's current likes and dislikes" (Moran 2012, 322). Might such demand be independent of one's like and dislikes but dependent on one's valuing, love, or attachment?

Perhaps the idea is that valuing, love, and attachment are constituted by second-order attitudes with regard to our likes and dislikes, wants, and inclinations. For example, Frankfurt says:

When we ...care about something, we go beyond wanting it. We want to go on wanting it, at least until the goal has been reached. ...[W]e feel it as a lapse on our part if we neglect the desire, and we are disposed to take steps to refresh the desire if it should tend to fade. The caring entails, in other words, a commitment to the desire. (Frankfurt 2006, 18-19)

On Frankfurt's view, our first-order likes and dislikes are accountable to our second-order desires and wants. But now the distance between what is beautiful, loved, or valued and what is merely agreeable seems to have shrunk. What seemed at first as a demand that is independent of the agent's attitudes, seems now to reflect a hierarchical relation within the agent's attitudes and desires.

Furthermore, even if attachment might be sufficiently fundamental to make sense of an agent's dismissal of his or her more superficial likings, it would not make sense of the direction of fit that we are seeking to explain. For, as Moran notes, it is a defining feature of the value we recognize that one may be measured by it. The direction of fit is from the subject to the object one values. On the attitude-based view, however, the direction of fit is from the subject to him- or herself.

Another worry about the attitude-based view is that it fails to explain our commitment to the object we value; it only explains our commitment to our first-order attitudes with regard to this object. For example, if valuing one's dancing career involves wanting to want to save the ballet company, then coming to terms with the fact that the ballet company will close constitutes a failure of one's commitment to one's first-order desire to save the company, not a failure of one's commitment to one's dancing career.

After all, the dancer's second-order desire to want to save the company can persist even if he no longer wants to save the company. Furthermore, the dancer's second-order desire to want to save the ballet company cannot explain why ceasing to have this second-order desire may strike him as a failure of his commitment to his dancing career. The second-order desire is the source of his commitment; it is not itself what he is committed to. If the second-order desire lapses, then so does his commitment.

Niko Kolodny (2011) and Susan Wolf (2010) voice another objection to the attitude-based view. The objection maintains that we may value objects that are not worthy of valuing. A person who cares deeply about counting blades of grass on the asylum's lawn does not have reason to be devoted to his pursuit. Perhaps another example may be found in Frankfurt's famous "willing addict," who affirms his irresistible urge to use drugs (Frankfurt 1988). The fact that such instances of valuing do not give rise to reasons of valuing is reflected in the fact that the rest of us lack minimal reasons not to disrupt such activities (Kolodny 2011, 74). So the mere fact that a person values something does not imply that he is reasonably devoted to it.

But Kolodny's most fundamental worry about the attitude-based account does not concern *what* gives us reasons to be devoted to an activity or pursuit, but the *way* in which it gives such reasons. Kolodny subscribes to the view that an object is a source of reasons if and only if there is a basic normative principle (a principle not derivable from further principles) according to which standing in a certain relation to the object gives us reason to bring it about, engage with it, honor it, or act in some way for the sake of it (Kolodny 2011, 49). Thus, if one's attitude gives one reason for action then the attitude

itself must be valuable, that is, it must be something that one has reason to bring about, or engage with, or honor, or act for the sake of. But, Kolodny argues, the attitude-based view does not hold that the attitudes that give us reason for valuing are themselves valuable. For example, on the attitude-based view I considered above, a second-order desire is the source of the dancer's commitment to dance without itself being an object of the dancer's commitment. Unlike attitude-based views, Kolodny wants to explain what gives us reason for valuing and devotion by appealing to a source of reasons that is itself of value.

Thus, as an alternative to the attitude-based view, Kolodny has offered a value-based view, according to which "a certain kind of attitude-based value is realized by one's pursuing an independently worthwhile aim that one cares about" (Kolodny 2011, 73). According to Kolodny, "[t]he fact that the dancer cares about dance means that his pursuing dance constitutes something of further value—a value that would not be constituted if he did not care about it... This further value is what provides him with reason to stave off a change" (Kolodny 2011, 73). Kolodny notes that Susan Wolf has invoked a similar attitude-constituted value (Kolodny 2011, 73). Wolf argues that meaning in life arises from finding fulfillment in one's involvement in objects that are valuable independently of oneself (Wolf 2010, 22). Our subjective experience is therefore crucial to finding meaning in life, but it reflects objective values and it is accountable to them. When we find our engagement in a certain project or activity fulfilling, we "find it such as to be characterizable in terms that would portray it as (objectively) good" (Wolf 2010, 24). Perhaps, then, reasons of valuing in general, and reasons of devotion in particular, are explained by the value of caring about an independently valuable object.

In considering this proposal, we should first note that neither Kolodny nor Wolf is suggesting that the goal of valuing is to give meaning to our lives or to realize the attitude-constituted value of valuing valuable things. Unlike those who hold that well-being is the only rational end to which all our attachments are instrumental, Kolodny and Wolf do not propose that valuing is merely instrumental to the value it constitutes. Rather, the suggestion is that since valuing valuable things is itself valuable, or since it gives meaning to our lives, we have reason to be devoted to the objects we value.

My worry about this proposal is that the value in question does not seem to play a justifying role in our devotion to the object we value. Our devotion looks outward, to its object, not inward, to its value as an attitude. Kolodny anticipates this worry. He insists that “this special value [the attitude-based value, O. N.] rises to deliberative salience ... where we step back from our relationships and pursuits, and survey our lives and their sources of meaning as a whole” (Kolodny 2011, 76). It is in moments of reflection on our life that we consider what we really care about and recognize the value of valuing something worthwhile.

Although Kolodny is surely right that in moments of reflection we step back and consider our history with the object we value and our feelings toward it, this does not yet show that in such moments we are concerned with the value of caring about a valuable object. To be sure, as we reflect on our past and on our feelings, we might come to realize that further devotion would not make sense or would not be justified, but I do not think that in our reflection on what we value we seek to find a further value constituted by our valuing it. What we seek is a better understanding of our reasons with regard to the

valued object, not an understanding of what further value they serve. Even in moments of reflection, taking the value of caring about valuable things as the ground of our devotion is having one thought too many.

Matters are slightly different when we think of the attitude-constituted value in terms of meaning. Finding that our relationship with someone, or that our devotion to a vocation, has given meaning to our lives does indeed seem relevant to our continued devotion to it. But, here too, I do not think that our reasons to be devoted to it are given by the meaning constituted by our valuing it; rather, it seems to me that the fact that valuing an object has given meaning to one's life *indicates* that one has strong reasons to be devoted to it. That is to say, the value constituted by valuing worthy objects depends on one's reasons to be devoted to such objects; it is not itself a reason to be devoted to them. If the dancer did not recognize reason to be devoted to his dancing career, his devotion to dance would not play a major role in making his life meaningful.

This objection is essentially the same as the one I raised with regard to the attitude-based account of devotion. Just as the attitude-based account conditions our reasons for devotion to an object on the fact that we value it, so the value-based account conditions our reasons for devotion on the fact that we value a valuable object. On both views, we have reason to be devoted to what we value, and to resist a change in what we value, only insofar as we value it. But, as I noted in considering the attitude-based view, our valuing is itself a response to the reasons we take ourselves to have with regard to the valued object. Since the fact that we value something worth valuing is partly explained by our singular reasons with regard to it, it cannot fully explain such reasons.

For similar reasons, Proust's protagonist comes to believe that if something like the attitude-based view is correct, then love is an illusion:

I said to myself sadly that this love of ours, in so far as it is a love for one particular creature, is not perhaps a very real thing, since, though associations of pleasant or painful musings can attach it for a time to a woman to the extent of making us believe that it has been inspired by her in a logically necessary way, if on the other hand we detach ourselves deliberately or unconsciously from those associations, this love, as though it were in fact spontaneous and sprang from ourselves alone, will revive in order to bestow itself on another woman. (Proust 1930/1992, 299)

Perhaps such a deflationary account of love and valuing is correct. John Gardner, too, conceives of our devotion to what we value as involving an illusion. Like Kolodny, Gardner argues that pursuing valuable goals is itself valuable (Gardner 2015, 14). But Gardner also believes that “from the inside, one does not think of the extra value as lying, content-independently, in the fact of commitment to the goal. To be committed one must think of the extra value as lying in the goal itself, meaning in its content” (Gardner 2015, 15). Gardner concludes that the extra value of our goals is a “rationally necessary illusion” (Gardner 2015, 15).

It seems to me that Gardner goes too far in holding that in order to value anything we must acquiesce in a necessary illusion according to which it is more valuable than other things. As I argued in the previous sections of this chapter, we may have reasons for valuing an object that others lack, and valuing consists in taking oneself to have (singular) reasons of valuing, particular to oneself, with regard to the object one values. Such reasons account for the varying commitments different individuals have with regard to one and the same object. But I agree with Gardner that being committed to an object we value occasionally involves disregarding our reasons *for* or *against* valuing it and

taking only our reasons *of* valuing into account. Our reasons *of* valuing explain why we often feel compelled to continue to value an object even when we may rationally forgo it.

The problem, I think, stems from the very notion of valuing. Proper valuing may often involve devotion that is unconditional and therefore oblivious to one's reasons for and against valuing. The projects, ideals, persons, activities, traditions, and goals that are dearest to us often call for conviction, trust, or faith. Sometimes, we should believe in a project even though the evidence strongly suggests that it will fail miserably; other times we should believe in a relationship even though everyone pleads with us to terminate it. Of course, there are also times when we should let go, give up the goal or withdraw from the relationship. But without commitment, without faith, we cannot value at all.

Consider an example of a different sort. In his speech on the steps of the State Capitol in Montgomery, Alabama, Martin Luther King famously said: "The arc of the moral universe is long, but it bends toward justice." MLK thereby expressed his commitment to justice and his faith that justice will prevail. The force of his exclamation turned, to a large degree, on the fact that there are many reasons to *doubt* its optimism. And yet faith in justice need not and should not depend on one's empirical judgment that justice is likely to prevail.

Devotion often requires a degree of obliviousness to the reasons for and against one's devotion. This may seem objectionable, and yet the persons and things we value most give meaning to our lives partly, but crucially because of our unwavering devotion to them. We therefore have reason to be devoted to what we value in ways that

occasionally block our view of our particular reasons for and against valuing and devotion.

It is not merely the case that the dancer has reason to stave off change, but that he believes that ceasing to care about dance, or no longer valuing it, would constitute a failure of his commitment to dance, a failure to respond to his reasons of valuing with regard to it. Indeed, the dancer's devotion to dance, just as Marcel's devotion to Gilberte, reflects his judgment that he should not give it up even if he has strong reason to do so. Even if the dancer correctly judges that he has most reason to let go and become a successful, fulfilled theater director who is engaged in something valuable, by following this path he will fail to respond to his reasons of devotion to dance.

Similarly, Marcel might move away from France and form other ties to people and places, thereby ceasing to be tormented by Gilberte and other attachments he has in the present, but such

new friendships for places and people are based upon forgetfulness of the old... my reason precisely thought that I could envisage without dread the prospect of a life in which I should be forever separated from people all memory of whom I should lose, and it was by way of consolation that my mind was offering to my heart a promise of oblivion which succeeded only in sharpening the edge of its despair. (Proust 1930/1992, 339; Moran 2012, 328)

Forming new attachments instead of past ones involves forgetting or becoming indifferent to our past commitments and reasons of devotion. Moreover, with the emergence of new objects of attachment, the objects we were once devoted to—objects that seemed utterly indispensable to us—are revealed as objects we were able and willing to do without. The fact that we forewent them casts doubt on our past devotion to them. Maybe we were able to forgo them all along? Maybe we were willing to give them up

even when we thought we could never live without them? If so, our devotion may be exposed for the conditional commitment it really was.

The very fact that we survived the loss of what we loved most may seem to retroactively undermine our love and devotion to it. The further thought that we might very well cease to mind our loss and value new objects in its stead only sharpens the edge of our despair, to use Proust's apt phrase. Such despair is fraught with self-alienation. When you lose your beloved or fail your greatest ambition, it is tempting to think that you are not the person you were when the object you valued was still within reach. And there is a sense in which this is true, insofar as your identity is understood in terms of the persons and things you love and value. But there is also a crucial sense in which you are exactly the same person you were before the loss occurred. If there is to be a subject who suffers loss, there must be a subject who survives it. It is because you are now the very same person you were then, that the loss is *your* loss.¹⁹

To alleviate the pain we experience when we suffer loss we may turn our gaze away—indeed, we may have most reason to turn our gaze away. Nevertheless, from the point of view of our commitment to the lost object turning our gaze would not help, for our reason to feel the pain would stay put. In other words, our pain is an expression of our

¹⁹ "...our dread of a future in which we must forgo the sight of faces and the sound of voices which we love and from which today we derive our dearest joy, this dread, far from being dissipated, is intensified, if to the pain of such a privation we feel that there will be added what seems to us now in anticipation more painful still: not to feel it as a pain at all—to remain indifferent; for then our old self would have changed, it would then be not merely the charm of our family, our mistress, our friends that had ceased to environ us, but our affection for them would have been so completely eradicated from our hearts, of which today it is so conspicuous an element, that we should be able to enjoy a life apart from them, the very thought of which today makes us recoil in horror; so that it would be in a real sense the death of the self, a death followed, it is true, by resurrection, but in a different self, to the love of which the elements of the old self that are condemned to die cannot bring themselves to aspire" (Proust 1930/1992, 338)

commitment, a futile yet necessary rebellion against the fact of loss. It is futile because there is no longer an object to be devoted *to*; it is necessary because there *was* an object to which we were devoted. Thus, devotion must give way at some point, but it must also go down fighting. I therefore do not mean to claim that we should never accept loss, or that we should never consider what justifies our devotion to the objects we value. Rather, my point is that devotion is normally resistant to its own rational basis and that sometimes the voice of reason should be silenced, while other times it should prevail.

Consider a project you value. It is of the nature of projects that they come to an end. If you pursue a project you value then you presumably wish that the project will come to a successful conclusion rather than end in failure. Let us assume that you complete your project successfully. Though you may continue to have reason to take pride in your success, you are likely to move on to new projects. Nevertheless, there are various ways in which even the successful completion of one's project can be painful. To see this, one only needs to consider the difficulty of retiring from a successful career that comes to its natural end, or watching one's children leave home as independent adults. It can be difficult to let go of the project or goal, but one *should* let go, for the attachment had run its course. Letting go of a project one has completed contrasts with ceasing to care about a project before or regardless of its completion. To be sure, one may have strong reason to quite the project and embark on a different one instead, and yet insofar as one properly values the project one should be committed to it. We may inappropriately hold on to a project that is no longer there, but excessive willingness to let go of a project may imply that we are not genuinely committed to it. Thus, we may lack reasons of valuing that we wish we still had and we may have reasons of valuing that we fail to

acknowledge. These cases are examples of pathologies of devotion. Devotion caves in on itself when it is permanently oblivious to the reasons that support it as well as when it constantly looks back on them.

There can be real tensions between the two points of view essential to valuing. On the one hand, our reasons *of* valuing and our commitment to what we value, and, on the other hand, our reasons *for* and *against* valuing and our acknowledgment of those reasons. The conflicts between these different kinds of reasons cannot be resolved by a better theory of valuing, or normativity, or whatever. Nor is it plausible that such tensions are the avoidable result of philosophical misunderstanding. Rather, these are conflicts and tensions that are essential to human life and experience, and that cannot be preempted or dissolved in advance but can only be confronted, instance-by-instance, where and when they arise. Whether we succeed or fail in our responses to such crises of meaning, the anguish and dread they involve cannot be averted, and Frost's question inevitably resounds:

Ah, when to the heart of man
Was it ever less than a treason
To go with the drift of things,
To yield with a grace to reason,
And bow and accept the end
Of a love or a season?

(Frost 1913)

Chapter III: Regretting and Being Oneself

Introduction

When asked to say who we are, we normally appeal to the persons we love and the things we value. In the previous chapters we considered the rationality of such attachments to persons and things. In this chapter, we turn to a darker side of our identities: our regrets. Regret is dark as in “wretched,” but it is also *in* the dark as in “hidden.” Our regrets are often the last things we disclose when we explain ourselves, but they are likely to be the most revealing. They express our yearning for persons and things that are no longer there to be loved or valued, and our longing for activities, actions, or choices, that can no longer be engaged with, performed, or made. As such, our regrets always aim at absences, though, at the very same time, they also aim at what is lamentably present: decisions, actions, events, and utterances that cannot be reversed, undone, or erased.

Initially, the notion of regret might seem quite straightforward and therefore unworthy of philosophical inquiry. It is natural to assume that regret is an expression of our negative judgment of past events, actions, or choices. As such, regret is merely an aftertaste of reason, so to speak, and any question that might be asked of it should be addressed to the judgment expressed by it. But on closer inspection, regret seems to normally and properly part ways with judgment. As Bernard Williams famously noted, mistakes or wrongs done in the past often lead to attachments that give meaning and worth to one’s life in the present (Williams 1981b). In such cases, regret might not be called for, or if it is called for, it may be remarkably difficult to summon. Alternatively,

our regrets might be the last stronghold of those persons and things we had lost or failed long ago. Even if losing them was overall justified or inevitable, or if we had good reason to fail them, these persons and things might continue to claim their due in the shadows of our minds.

The decoupling of regret and judgment might seem to suggest that regret is not liable to rational assessment, but should be viewed instead as a merely psychological phenomenon. The strains of regret, and its insistence on that which cannot be realistically hoped for, may further lead us to conclude that we should do away with this psychological hang-up, insofar as we can. I argue against this line of thought in the first section of this chapter. There are, I will claim, genuine reasons for regret, even if we often have strong reason not to want to regret, or not to dwell on the past.

Each of the subsequent three sections is dedicated to a different kind of regret: regretting a mistake (section 2), regretting a loss (section 3), and regretting an attachment (section 4). The second section explains why we may lack reason to regret a past mistake even if it remains unjustified in light of all relevant information. In the third section, I distinguish between regretting a decision or choice we made, and regretting the loss it involved. I argue that we often have reason to regret loss even if we do not have reason to regret the justified choice that brought it about. Both when we lack reason to regret a past mistake and when we have reason to regret a loss, our attitudes are rationally explained by the justified attachments we have in the present.

The fourth section of this chapter addresses the rationality of regretting our attachments themselves. I focus on political contexts, including a particular political

context that is of special importance to me. In 1948, immediately after Israel declared its independence, a war broke out between Israel and the surrounding Arab countries. Israel considers this war its war of independence, but during this war an estimated 700,000 Palestinians fled or were expelled from their homes (Morris 2003, 588). These refugees and their descendants number several million people today, divided between Jordan (2 million), Lebanon (427,000), Syria (477,700), the West Bank (788,100) and the Gaza Strip (1.1 million), with at least another quarter of a million internally displaced Palestinians in Israel. The displacement, dispossession, and dispersal of the Palestinian people is known to them as *an-Nakba*, meaning “catastrophe” or “disaster.” The Nakba Day is commemorated on May 15, the day after the Gregorian calendar date for Israeli Independence Day.

In recent years, there have been voices in Israel calling to forbid the commemoration of the Nakba. A law has been proposed to incarcerate for a period of three years anyone who marks Israeli Independence Day as a day of mourning. Eventually, a more moderate version of the law was passed, which grants the finance minister the power to reduce the budget of state-funded bodies that openly reject Israel as a Jewish and democratic state, or that mark the Nakba (Khoury 2012).

Many Israelis feel that the commemoration of the Nakba is akin to a rejection of Israel’s right to exist. Thus, they believe, an acknowledgment of the Palestinian disaster would be tantamount to a rejection of everything they hold dear. In the eyes of many Jewish Israelis, their relationships and projects, their culture and values, are all embedded in this event. If it were not for this “disaster,” they say, they would not have existed.

Thus, it is not uncommon to hear Israelis proclaim that if Israel were to acknowledge the Nakba, or to allow the Nakba to be widely acknowledged, this will mark the beginning of the end of the State of Israel.

Part of what I hope to do in this chapter is to offer a general account of regret that may help address this issue. That is to say, eventually I would like to consider whether my own acknowledgment of the wrongs committed during Israel's war of independence rationally undermines my attachment to Israel, its culture, its people, the values it stands for, and its institutions. Furthermore, I would like to consider whether, insofar as Israel was and is necessary for many personal attachments of those who, like me, see Israel as their home, the acknowledgment of the Nakba undermines these personal attachments? I will not consider these questions in depth here, as there are many related questions—about historical events, but also about the relevant object of regret, rejection, or affirmation—that must be settled before we can address the matter adequately. But I hope that what I say about regret may provide a framework for thinking about these questions, and I will briefly comment on them in the fourth section of the chapter. For now, let us move away from history, war, and politics, and consider whether we have reasons to regret anything at all.

1. Not Regretting

To regret an action or event is to wish it was not part of the actual past. This understanding of regret is not uncontroversial, but it is sufficiently plausible to fix what I take regret to mean in this chapter. One reason the above explication of regret might seem inadequate is that regret is often understood as essentially agential, that is, as referring to

the agent's own actions and choices, not to actions, choices, or events independent of the agent. But it seems to me that we can regret the circumstances in which we found ourselves without regretting the choices we made in those circumstances, and we can also regret actions of others, or misfortunes in which we were not involved. To be sure, agential forms of regret have distinctive features, but they share with non-agential forms of regret the wish that things were otherwise. In the next section I will consider whether a person is required to regret his or her own mistake and will trace this requirement to a person's agency in the present, but I will generally use the word "regret" more broadly to refer to a person's wishes with regard to the past.

There is an apparent reason to doubt that there are genuine normative reasons for regret. Unlike wishing for something in the future, wishing that the past were different is wishing for something that cannot be. Our regrets are always, and necessarily, frustrated. This is why regret is associated with discontent and alienation: when we regret the past we set ourselves against it, we wish our world were different than it actually is. Such discontent and alienation are even more acute in cases of remorse or repentance. The futility of such negative attitudes toward the past has led some to rule regret unreasonable. But the idea that regret is futile presupposes a certain understanding of practical reason, of its role and structure. I therefore begin by considering the role of the past in practical reasoning.

What should we do about the past? The question may seem odd. As we reason about what to do, we normally reason about the present or the future. Indeed, what else *can* we be reasoning about? Practical reason guides us, tells us what to do next, and as

such it is essentially forward-looking. The answer to our question might simply be “nothing,” for there is nothing we *can* do about the past. The past is beyond our control and therefore outside the jurisdiction of practical reason.

But the past is normally relevant to what we should do. It is relevant in the sense that it informs our present and future actions. For instance, when I have reason to apologize to you, the reason is given by an action I performed in the past, but offering my apology is what I should do now, or next, or in the future. And though sometimes people think of apologies or punishments, reparations or compensations, as ways of cancelling past actions, such cancellations cannot undo the very fact that the actions they defuse took place. At most—and this is also questionable—they can undo the wrong, harm, or cost of past actions. So even though we cannot do anything *in* the past, the past often bears on what we should do in the future or in the present. In this sense we may indeed ask what we should do (now or in the future) about the past.²⁰

In fact, we can ask an even broader question. Practical reason does not only guide action, it also guides attitude. Admittedly, unlike our actions, our attitudes are normally not voluntary, and yet we do cite reasons for or against attitudes. Judgment-sensitive attitudes are attitudes that are responsive to our judgments about reasons (Scanlon 1998, 19-21). My fear of your dog reflects or commits me to the judgment that your dog is dangerous. If you persuade me that the dog is harmless then, insofar as my fear is rational

²⁰ Note that the past may be relevant to future consequences, but it may also be relevant to what we should do now or in the future *regardless of the consequences*. What I did to offend you might explain why an apology will maximize utility, or value, or whatever, but it might also give me reason to apologize independently of whether doing so will maximize anything. The conception of practical reason as essentially about the future does not entail a consequentialist conception of practical reason.

it would dissipate. Not only our actions, but also our judgment-sensitive attitudes may be about the past. My anger at my friend may reflect my judgment that she betrayed my trust. So in addition to asking what we should do about the past we may ask what attitudes should we take with regard to the past.

Still, one might insist that I have only shown that judgments about the past *bear* on the present and future actions and attitudes we should take; I have not shown that the past may be an appropriate *object* of our present and future attitudes and actions. After all, my friend's betrayal may justify my present anger *toward her*, not my present anger *toward the past betrayal*. Similarly, the fact that your dog attacked a passerby last week may justify my present fear *of your dog*, not my present fear *of the past attack*. The fact that the past is relevant to our attitudes and actions does not show that we may rationally adopt attitudes *about* or *toward* the past.

But it seems difficult to deny that I may reasonably be glad that something wonderful happened to me yesterday. To be sure, I may be glad because of the future consequences of what happened to me, but that need not be the case. My gladness need not be directed toward any person or thing in the present or future; it may be directed toward the beautiful music that I heard yesterday, or the success I had, or the kindness I encountered. The past may be celebrated or mourned regardless of the future; it may be what our attitudes and actions are about.

But the suspicion that the past is irrelevant might persist. On a broadly consequentialist view, our present and future attitudes and actions should be those that will lead to the overall best consequences. On such a view, the past is only relevant to our

attempts to bring about most value in the future. Insofar as the past may be an appropriate object of our attitudes and actions, it is merely instrumentally valuable for the sake of the single, ultimate end that lies ahead. Moreover, a consequentialist might subscribe to an experientialist account of value, on which our attitudes and actions should maximize positive experiences. It might then be thought that we should, all else being equal, prefer to have good, pleasant, or positive attitudes, as opposed to bad, painful, or negative ones. So, we might conclude, with regard to any action or event in the past, we should take a good, pleasant, or positive attitude rather than a bad, painful, or negative one, as long as doing so does not result in even worse consequences.

This line of thought may lead us to rule out regret, remorse, shame, guilt, a sense of loss, and other negative attitudes with regard to the past, for these are attitudes that, according to the reasoning we are considering, cannot be justified when taken on their own, that is, independently of good or valuable consequences they might give rise to. For example, Rüdiger Bittner has argued that it is never reasonable to regret what one did, for once we judge what we did as wrong or mistaken and take action to address it—e.g., apologize, compensate, take measures to avoid the recurrence of the action, etc.—then there is no point to feeling bad about the past action and there is reason not to (Bittner 1992). Bittner’s claim is in line with a tradition of thought that condemns regret and its host of anguished attitudes toward the past. Emerson, for instance, urges us to “[r]egret calamities if you can thereby help the sufferer; if not, attend to your own work” (Emerson 1841/2010, 45). And Nietzsche prescribed: “Never give room to repentance, but tell your-self at once: this would mean to add a second piece of stupidity to the first” (Nietzsche 1986, 323). The same idea is encapsulated in common proverbs, such as

“What’s done is done” and “No use in crying over spilled milk.” The painful feelings of regret seem redundant, even objectionable, once separated from their beneficial upshots and the truthful evaluation of the past action or event, for there seems to be no value in mere anguish. This is the *anti-regret view*.

It seems to me that the anti-regret view is often motivated by a worry about self-indulgence. When I ask whether or not I should regret my past action, I seem to be concerned with myself and my attitudes rather than with my past action and the steps I should take in the present and future in order to address it. My regret would not reverse the past action, nor would it cancel or mitigate what is wrong or objectionable about it. One might propose that by focusing on whether we should feel bad for what we did, we fail to take responsibility for it and replace the lessons we should draw for the future and the responsibility we should take for the past, with an unpleasant feeling or experience. We throw our hands in the air and wish for our pain to do our work for us, to set things right, to absolve us. Thus, proponents of the anti-regret view urge us to focus on what is valuable and forgo a preoccupation with our feelings and attitudes about past wrongs and misfortunes.

The anti-regret theorists are surely right that there comes a point when one should let go and move on rather than dwell on the past. But the claim that our attitudes toward the past should be influenced by forward-looking considerations leaves room for the idea that such attitudes should also be influenced by backward-looking considerations. We can take responsibility for what we did, address it, and correct our ways, and still regret that we did it. In other words, if the point of the anti-regret view is merely that we shouldn’t

dwell on the past, then this point is a long way from the claim that we should ignore the past altogether.

But if we understand the anti-regret view as stated, then I believe that it is mistaken. Not only does it not focus our attention on what is important, it directs our attention away from it. The anti-regret view has two main premises. First, it assumes that our attitudes in the present or future are properly guided by what attitudes would bring about the best consequences overall. This is the *consequentialist assumption*. Second, the anti-regret view assumes that regret is a negative experience or sensation that may be considered independently of the judgment that invokes it. This is the *experientialist assumption*. In light of these assumptions, it would then seem that, other things being equal, we should always be happy, or pleased, or joyful with regard to any action or event, unless taking such attitudes would lead to worse results overall. And so even with regard to horrible deeds we performed in the past (or horrible disasters that are approaching), it is not only that we should not regret or bemoan them; rather, we should, ideally, rejoice in them, for happiness is better, more pleasant and therefore more valuable, than pain, anguish, or indifference, and we should take on the most valuable attitudes. In other words, the anti-regret view commits its proponents to holding that when happiness or content are inappropriate it is only because they may lead to bad consequences and not in any way due to their object.

This is a strikingly implausible upshot. Regret cannot be guided by forward-looking considerations alone because it is essentially a cognitive attitude that reflects the agent's judgment of the past. A person who rejoices in a horrible action that she

acknowledges as such seems to manifest a deep incoherence, and it is only slightly less jarring to remain indifferent to one's own harm or good fortune. Such disconnect between judgment and attitude seems to manifest a rational failure. For insofar as a person is happy due to her belief that happiness is better, her happiness is not primarily concerned with the object that it is purportedly about, but with its own value as an attitude. It is in this sense that the anti-regret view leads us away from what is important: it holds that our attitudes are justified by their own value, or by their contribution to the most valuable consequences, rather than by their objects. But insofar as a person's attitude is indeed *about* an action or event, it cannot entirely rely on its own value as an attitude or on the value of its consequences, but must express the person's assessment of the action or event that is its object.

Thus, by reducing all practical questions to questions about present and future value we seem to posit present and future value as the underlying object of all our actions and attitudes. If practical reason is to guide actions and attitudes that are not about present or future value, then such actions and attitudes must be, at least partly, rationally determined by facts about their objects. I say "at least partly," because, as I said, we may often have most reason to let go of the past, to move on, and in such cases we should not regret, even though we have (pro-tanto) reason to regret. Thus, even if it is true that I should prefer to be happy, pleased, joyful, etc., for such attitudes are more valuable, or better, the potential objects of these attitudes might give me reason not to have such positive attitudes.

All this is to say that the anti-regret view fails to recognize that reasons for regret are normally given by the objects of regret. It is generally agreed that sometimes there are reasons that seem to count in favor of an attitude while being inappropriate as the rational basis for the attitude, that is, while being *the wrong kind of reasons* for the attitude. The phenomenon is often illustrated with regard to beliefs. Perhaps believing that I will do well in the exam will increase my chances of doing well in the exam, and I would therefore have reason *to want* to believe that I'll do well in the exam. But this is compatible with having conclusive reason *to believe* that I will not do well in the exam—e.g., perhaps I haven't studied for the exam at all (Hieronymi 2005; Hieronymi 2013). The same phenomenon might occur in the case of intention. The fact that I will receive a great amount of money if I intend to drink the toxin might be a reason to want to intend to drink the toxin, but it would not be an appropriate reason to intend to drink the poison. Perhaps it is not even possible to intend to drink the poison for the sake of having the intention to drink the poison.

Similarly, even if we have strong reasons not to want to regret past events, and even if, as some argue (Schroeder 2012), these reasons can amount to decisive overall reason not to regret—as when we have most reason not to dwell on the past—there is conceptual room for pro-tanto reasons to regret. Reasons to regret are not primarily given by facts about what attitude would be most valuable or good to have, but by facts about the object of regret, that is, the regrettable action or event. The main problem with the anti-regret view is not that it offers the wrong kind of reasons against regret, but that it fails even to recognize the distinction between reasons given by the object of regret and reasons given by the attitude of regret. In other words, the view fails to distinguish

between object-given and state-given reasons for regret (Parfit 1984, 169; Parfit 2011, 420-432). Once the distinction is in view, we may ask: what are object-given reasons for regret?

To see more clearly the difference between reasons for and against regret and reasons for and against wanting to regret, consider the following, rather odd predicament. You are at the mercy of an anti-regret tyrant. The tyrant tells you that unless you cease to regret the one choice you regret most in life, he will cause you great anguish. The tyrant further clarifies that he is not asking you to change your mind about whether you should have made the regrettable choice; he only demands that you will no longer be bothered by the fact that you made the regrettable choice. By threatening you, the tyrant gives you a strong reason to want not to regret the choice.

But now a sophisticated anti-regret tyrant strolls by and, as he notices the situation, steps up and offers his advice to the first tyrant. "Threats are redundant," the sophisticated tyrant says to the first tyrant, "you should simply point out to the prisoner that by ceasing to regret so intensely she will be relieved of great anguish." The sophisticated tyrant realizes that you already have strong reason not to want to regret, namely, that regretting so intensely involves great anguish. The sophisticated tyrant thereby joins Bittner and the anti-regret theorists in holding that by regretting the past action you are failing to respond appropriately to your reasons.

But in response to the sophisticated tyrant you are unlikely to smack your forehead and wonder how have you missed such a relevant consideration against regret. For your regret about the past choice is not due to a failure to notice the great anguish that

regretting the choice involves, nor is it due to your failure to appreciate your reasons to want not to experience anguish. On the other hand, not every past action that you believe you should not have performed is an action you regret and are anguished by. And just as your regret is not due to a failure to appreciate the anguish that regret involves, your lack of regret is normally not due to your appreciation of the anguish that regret would involve.

The desirability of not regretting is external to your primary reasons for regret, which are given by facts about the object of your regret rather than by facts about the present or future value of the attitude of regret. The attitudes we should take toward the past may be influenced by facts about future or present value, but they are primarily determined by facts about the past. If I have done something horrible, something that I absolutely should not have done, then I have strong reason to regret it, regardless of whether there is anything good in regretting, or whether any good will come of it. To rejoice in the past wrong I had committed simply because such attitude would be more valuable all things considered is to set myself against my own assessment of the wrong I had committed. Similarly, to feel anguish about a past event that I judge favorably involves a rational failure not because of the badness of anguish, but because, according to my own judgment of the event, I lack reason to feel anguish.

To be sure, the worry of self-indulgence, which, I have suggested, underlies the anti-regret view, is a real one. In the Israeli left there is a term for a particular manifestation of such self-indulgence: “shooting and crying.” The term applies to Israeli soldiers who claim to regret their contributions to past wrongs while continuing to

contribute to new wrongs of the very same nature. The soldiers' willingness to continue to commit wrongs they claim to regret renders their regret void: they are "crying" over wrongs they are continuing to commit. To be sure, the soldiers' tears might reflect their regret over the circumstances that required their engagement in the armed conflict. In this case their continued engagement in the conflict would not void their regret, for the regrettable circumstances that require their engagement might still be in place. Such regret would not take the soldiers' own actions and choices as its object. Indeed, on this interpretation, the soldiers' tears absolve them, for the tears imply that the soldiers are doing the right thing in the regrettable circumstances, which they recognize as such. Thus, if the soldiers were to genuinely regret their own actions, and believe they have acted wrongly, such regret should reflect a real change in their views and actions. The fact that we may have reason to regret wrongs or mistakes does not imply that we may continue to commit the same wrongs or mistakes as long as we properly regret them. Nor does it imply that regretting is *all* that the past requires of us. But we should also be wary of holding that our attitudes should be guided solely by forward-looking considerations. An attitude toward the past that is only concerned with the future is as misguided as an attitude toward the past that is utterly indifferent to the future.

2. Regretting a Mistake

Once we acknowledge that reasons for regret are distinct from reasons to want to regret and from considerations of present and future value, we may want to know what they are. That is, what facts about the past call for regret in the first place? A natural and compelling answer to this question is that a person has reason to regret his or her past mistakes. In this section I consider grave mistakes that we have most reason *not* to regret,

and even reason to affirm them, or to be glad we made them. In the next section I consider cases in which we have reason to regret *justified* choices and actions, whether our own or others', as well as events out of our control. In each of the next two sections we will see that whether one has reason for regret is not only a function of the quality of one's past choices, but also of the justified attachments one has in the present. We will thereby arrive at a more nuanced understanding of the rationality of regret. One does not regret the past from the timeless point of view of judgment, the point of view from which one sees what one ought to have done; rather, one finds reasons for regret when one looks at the past from the present, from the point of view of the person one is today.

*

Sometimes we may reasonably be glad that we made a grave mistake. I will consider two cases that are discussed in the literature. First, there is the case of the young girl's child, originally proposed by Derek Parfit (Parfit 1984, 357-361) and brought here with some adjustments. Consider the decision of a fourteen-year-old girl to conceive a child. There may be strong objections to the girl's decision. Not only would it be wrong to bring a child into the world when the girl is not sufficiently mature to properly attend to the child's upbringing, but by having a child at fourteen the girl would also deny herself of various opportunities and experiences that she has strong reason to want. Let us assume therefore that the girl has most reason to wait a few years before becoming a mother. The question is, if the girl has a child at fourteen contrary to what she has most reason to do, does she later have reason to regret her decision and to wish she had not given birth to her child?

As an aside, it is noteworthy that the young mother's predicament is quite common even amongst those who had children later in life. It is not unusual to hear a person honestly express regret with regard to his or her choice of partner, with regard to the whole relationship, until the moment we remind him or her that without their partner and relationship they would not have had their children. And so it might be that the thing a person regrets most about her life is her choice of partner, and the thing that she celebrates most is her child. A person in such predicament faces the same problem as the young mother. Such a person judges that she should not have made a decision that was necessary for the birth of her children.

Second, consider parents to a deaf infant who are faced with the option of curing the infant's deafness. The parents may decide to allow a medical procedure that gives the infant a cochlear implant that would enable him to function normally in the hearing community (Harman 2009, 177-178). Installing the implant involves an invasive medical surgery, so there may be significant reasons against it. Still, let us assume, the parents have most reason to approve the medical procedure, given that, arguably, a life as a hearing person would be overall better for the child. Suppose, however, that the parents do not approve the procedure and accept the infant's deafness, and that the infant grows up to be a happy, fulfilled deaf adult. Furthermore, the child's deafness leads him to have valuable friendships, projects, and experiences that he would not have had as a hearing person. He becomes part of the deaf community, and sign language becomes a form of expression that he values in and of itself. Given that the parents had most reason to approve the medical procedure and cure the infant's deafness, do the parents now have

most reason to regret their decision, that is, to wish their child had a very different life, the life of a hearing person?

The reason these cases are controversial is that we may want to hold that the agents involved have overall most reason *not* to regret their past decisions, and that it is reasonable for them to be glad they made the choices they actually made. The young mother loves her child and seems to have most reason to affirm the child's existence despite the fact that her choice to conceive was objectionable; the parents love their deaf child as he actually is, and therefore seem to have most reason to celebrate rather than regret the mistaken decision that made him the person he is today. In both cases, if the subsequent affirmation is reasonable, then it may seem to warrant a change in our evaluation of the past choice: the mother might insist that it was not a mistake to conceive at a young age and the parents may hold that it was not a mistake to refuse the implant. But does the fact that the agents may affirm their choices in the present—that is, prefer to have made the choice they actually made—really entail that the choices were not mistaken? Alternatively, does the fact that the choices were mistaken really imply that the agents ought to regret them?

It is tempting to assume, quite simply, that we should regret what we should not have done. That is, if an action or choice is wrong, objectionable, unjustified, or mistaken, then we should regret it. Take Elizabeth Harman's example of "I'll be glad I did it" reasoning:

Last night I was trying to decide whether to work on this paper or go out to a movie. I realized that if I worked on this paper, then today I would be glad I did it. This enabled me to see that I should work on the paper rather than going out to a movie. (Harman 2009, 177)

The fact that Harman has most reason to work on the paper explains why, if she does work on the paper, then tomorrow she'll have reason to be glad she did (Harman 2009, 194). The same reasoning may explain why she has reason to wish she did not go out to a movie: she had most reason not to do so. That is why we can engage in "I'll regret it" reasoning, too. So, it might be thought, in general, we have reason to regret not doing what we had most reason to do, what we ought to have done.

On this view, if the agents in the two cases—the young mother and the parents of the deaf child—made mistaken choices, then they should regret them, that is, they should wish they had made the right choices, the choices they had most reason to make. The view implies that we should reject our initial judgment that the young mother may reasonably affirm her decision to conceive and that the parents may reasonably affirm their decision not to give their child the implant. By attributing this view of regret to the mother and the parents we can explain why they might feel rational pressure to defend their past decisions. Given their love for their children, regret may seem utterly unreasonable to these agents, and so if it is true that one should regret one's mistakes, as the view under consideration maintains, then in order to justify their lack of regret the mother and parents would have to justify their past decisions and thereby undermine the apparent call for regret.

We may try to resolve the conundrum by appealing to the difference between what an agent has most reason to do in light of the information available to him or her, and what the agent has most reason to do in light of all the relevant facts. In other words, we may distinguish between *subjective justification* and *objective justification*. It is often

true that in retrospect we learn facts that are relevant to what we ought to have done but that we were reasonably unaware of at the time of choice. In such cases, we may regret a choice that we were subjectively justified in making due to the fact that we were objectively unjustified in making it. Other times, we may make a decision that is mistaken in light of what we know and yet it turns out to be the correct decision. Suppose, for example, that you do not take an umbrella to work despite learning from the forecast that it is very likely to rain and therefore having most reason to take an umbrella to work. Your choice, we may assume, is *subjectively unjustified*. Nevertheless, as it turns out, it does not rain. Your choice, we might then say, was *objectively justified*. You may therefore be glad that you made the subjectively unjustified choice of leaving your umbrella at home. Thus, we may reasonably affirm subjectively unjustified choices if they turn out to be objectively justified.

But the distinction between subjective and objective justification is not helpful in the two cases we are considering because the choices in question are *objectively unjustified*. We are assuming that in light of all the relevant facts, the young mother should not have had a child at fourteen and the parents should have approved the medical procedure that would have enabled their child to hear. Neither the mother nor the parents learn new facts that they did not know at the time of choice. For instance, the mother could have known that she will become attached to her child and will therefore be reluctant to regret her choice to conceive. But this information does not undermine her decisive reason not to conceive. Since in both the mother's case and the case of the parents, lack of regret seems warranted even if there is no gap between what the agents

were subjectively required to do and what they were objectively required to do, the distinction cannot help us in addressing these cases.

Another possible response to these cases holds that the mother and the parents may be reasonably ambivalent about their past actions and choices. They may regret their past choices insofar as they recognize that these choices were mistaken, but at the same time they may affirm their past choices insofar as they love their children as they actually are. R. Jay Wallace describes such ambivalence as “having regrets,” and distinguishes it from a stable, overall preference that things should have been otherwise, which he calls “all-in regret” (Wallace 2013, 45-46, 51).

But we may worry that the mother’s and the parents’ love and commitment to their children calls for something stronger, that is, for a definite affirmation of the choices that led to their children’s existence as the individuals they actually are. This is an instance of what Wallace calls “unconditional affirmation”:

If we are attached to an individual or to a project, then we will typically affirm the direct objects of our affirmation in a distinctively unconditional way; this in turn commits us to affirming their necessary constitutive and historical and normative conditions in a way that is similarly unconditional, and precludes our regretting that those conditions obtained. (Wallace 2013, 77)

Given that the mother’s and the parents’ past choices were unjustified, and that they should have chosen differently, may they now reasonably prefer overall to have chosen as they in fact did?

Furthermore, even if we put aside the requirement of unconditional affirmation, it might seem that the fact that the mother and the parents did what they had most reason *not* to do, implies that they should prefer overall not to have done so, no matter how

attached they are to the consequences of those actions or choices. In other words, past unjustified choices, such as those made by the mother and the parents, may seem to call for all-in regret rather than mere ambivalence. Thus, the problem with the current answer to our puzzle is twofold. First, ambivalence may often seem incompatible with our reason for unconditional affirmation of past mistakes, and, second, ambivalence may often seem incompatible with our decisive reason to (all-in) regret past mistakes.

However, it has been recently argued that, on some occasions, agents may reasonably *not* regret, and even affirm, their mistakes and unjustified choices. Harman has argued this much about both of the above cases:

Sometimes it is (or will be) reasonable to prefer an outcome even though the alternative would have been better (in all the ways one should care about). It is reasonable of parents to prefer that their adult deaf child have come to be who she is, even though it would have been better (in all the ways they should care about) if their child had been cured of deafness. A teenager who has chosen to conceive will later be reasonable in preferring that her child exists, even though it would be better (in all the ways she should care about at the time she chooses) if she waits to conceive later. (Harman 2009, 188)

Harman describes the cases in evaluative terms and argues that one may reasonably prefer a worse outcome to a better one. But Harman's solution underestimates what is at stake. The issue is not whether we should always prefer the best option. There may be cases in which past obligations commit us to certain sub-optimal alternatives or where one may reasonably choose between sufficiently worthy options without being rationally required to choose the best option. What is striking about the cases we are considering is that if a lack of regret is reasonable then one may prefer an option one had most reason not to prefer at the time the choice was made. It is not that the option one prefers was less-than-best; it is that the option one prefers was wrong, mistaken, or unjustified.

Harman suggests that deaf activists who call upon parents not to cure their babies of deafness wrongly assume that “their preferences for their actual lives are reasonable if and only if their actual lives are best for them” (ibid.) Harman believes that if deaf activists realize that they may reasonably prefer their lives to the better lives they would have had, they could acknowledge that parents should allow the implant that offers their babies better lives. But deaf activists might continue to worry that if their parents should have cured them of deafness, then their parents should regret not doing so, and they, the deaf activists themselves, should regret that their parents did not cure them of deafness. The question at stake is not what would be best for the child, but what the parents should do.

Indeed, if Harman is right that the relevant mistake is to assume that one should always prefer the best option, then the reasonableness of the parents’ later preference for their child’s actual life as a deaf person seems to entail that the original choice not to give him the implant was also reasonable. In other words, it is only in light of the fact that it was reasonable to prefer a sub-optimal option to begin with that in retrospect it is reasonable not to regret the choice. But this is not the solution we are looking for, nor is it the solution Harman is looking for. Harman wants to say that it *is* a mistake not to cure the child’s deafness. She believes that it would have been better for the child to be able to hear and that in this case the parents *should* choose what is best for their child. But Harman also believes that parents who did not approve the implant may reasonably affirm this mistake in retrospect (Harman 2009, 189). So Harman still faces the problem: how might it be reasonable of the parents to prefer their mistake to what they should have done?

The problematic assumption is not that we should always prefer the best—an assumption that has been doubted by many philosophers—but that we should always prefer to have done what we had most reason to do—a much less controversial assumption. The latter, more plausible assumption is not one that may be easily dismissed, not even in light of our considered judgments about the two cases. And even if the two cases are seen as counterexamples to the assumption that one should prefer to have done what one had most reason to do, we must still explain how is it that this plausible assumption does not hold in these cases.

Recently, R. Jay Wallace has offered an account of regret that aims to do justice to our considered judgments about the two cases. In considering the case of the young girl's child, Wallace says that two desiderata must be met. First, we must account for the fact that the mother should affirm her child's existence, that is, she should not regret conceiving at a young age but rather be glad that she did. Second, we must respect the fact that the considerations that counted against conceiving do not alter after the child is born. That is to say, the existence of the child does not call for a change in our judgment about the girl's decision to conceive at fourteen: the decision remains unjustified after the child is born.

To explain these two considered judgments, Wallace proposes that we abandon

the assumption that our prospective and retrospective attitudes should be understood in evaluative terms, as involving (for instance) all-things-considered assessments of the child's existence. Instead, I propose that we think of the case normatively, attending to the changes in the deliberative situations of agents that are induced by the birth of the young girl's child. (Wallace 2013, 88-89)

Once the child is born, says Wallace, the girl has reasons she did not have before, reasons having to do with her actual child, the individual human being who stands in a significant

relationship of attachment to her. Wallace claims that the fact that “it is no longer an open question whether to have a child or not ... alters the normative landscape that [the young girl] inhabits” (Wallace 2013, 89). So even if the young girl now has most reason not to regret her decision to conceive, she may still believe that “someone in the situation she earlier faced ought to put off motherhood until becoming more mature” (Wallace 2013, 90).

In other words, before deciding to have a child, the girl can recognize the reasons of attachment she will have once the child is born, but those reasons are not in place yet. Once she has the child, these reasons come to life, so to speak, and the girl has reasons she did not have before to affirm her choice to have a baby. The idea is that between the moment the girl makes the choice and the moment in which she reflects on it there is a significant change: the girl’s child comes into existence. This change involves a change in the girl’s “deliberative situation,” as Wallace calls it, which explains her compelling reason not to regret her choice.

But Wallace’s explanation does not address the apparent tension between the fact that the choice was mistaken and the fact that the girl should not regret it. The girl’s new deliberative situation—her reasons to affirm her choice—does not disarm the powerful assumption that one ought to prefer to have done what one had most reason to do. To insist that in the case of the young girl’s child this principle does not hold seems ad-hoc unless we offer some further explanation of why the principle does not hold.

Suppose I know that I will have reason to speak Italian once I arrive in Rome, six weeks from now (Nagel 1970, 63). I therefore have reason to start learning Italian now in

order to be able to speak Italian when I am in Italy. Suppose, further, that I do not learn Italian and arrive in Rome unable to speak the language. It is no longer an open question whether I learn Italian before arriving in Rome, for I have already failed to do so. My normative landscape and deliberative situation have changed: I now have new, compelling reason to make the best out of my visit given the fact that I do not speak Italian. In particular, I have reason to do my best to learn Italian while I am in Rome and not to agonize over my failure to learn Italian before my arrival. Still, it seems that regret would not be unreasonable in this situation, and that I should prefer to have studied Italian before my arrival. In particular, in the absence of new relevant information, it would be unreasonable to be glad that I did not study Italian. In order to explain the case of the young girl's child we must explain how it differs from the case of the tourist in Rome. In both cases the past mistake leads to a new deliberative situation that could have been anticipated in advance and that the agent had most reason to avoid; in both cases the agent's new deliberative situation gives him or her new reasons not to regret but rather affirm their mistake. Why is it that in one case it is reasonable for the agent to affirm her past choice and in the other case it is not?

Wallace calls our attention to the fact that sometimes objectionable, mistaken, or unjustified choices may lead to new attachments that give rise to compelling reason not to regret those choices. This is a valuable and relevant point. But what remains perplexing is that such reasons may cancel, silence, or undercut our reasons for regret. After all, we already knew that we have strong reason not to want to regret regrettable choices. When considering the anti-regret view, we saw that the fact that deep regret is a great burden seems to give us compelling reason to want not to regret our mistakes. However, such

considerations do not cancel, silence, or undercut our reasons for regret. Even if such considerations give us reason not to dwell on our past mistakes, we cannot reasonably prefer to have made those mistakes. Thus, the fact that I want to make the best out of my visit to Rome may give me reason not to dwell on my failure to study Italian, but I cannot reasonably affirm my failure to study Italian.

What is troubling in the case of the young mother and the case of the parents of a deaf infant is that the new reasons to affirm their past choices, to which Wallace draws our attention, seem like reasons to *want* to affirm the past choices but that do not themselves justify such affirmation. The problem is that the mother's and the parents' reasons compel them to affirm a choice—to prefer to have made it—independently of whether the choice was justified.

We have seen that in order to make sense of the two cases it is not enough to insist that the agents may reasonably affirm their mistaken choices; we must explain *why* they may reasonably affirm their mistaken choices. And we have also seen that the answer to this question cannot be that it is sometimes reasonable to prefer sub-optimal options, nor is the answer that the agents have new reasons that they did not have at the time the choice was made. In order to proceed, we should consider the principle at work in “easy cases,” where a mistaken choice or decision should not be affirmed in retrospect. The principle underlying the easy cases would be the principle that misfires in the cases of the young mother's child and the deaf infant, assuming affirmation is reasonable in such cases. Once we articulate this principle, we will be in a better position to consider why it is not in force in the cases of young mother and the deaf infant.

Here are some easy cases, already mentioned above. When I have most reason to stay home and write but act contrary to what I have reason to do and go out to a movie, then I should regret my choice, or at least not affirm it. That is to say, I should, in retrospect, recognize that I acted contrary to what I had most reason to do and wish, or prefer, that I did not. In particular, I cannot reasonably prefer overall to have made the mistake. The same is true in the case of the tourist in Rome. The tourist has reason to regret not studying Italian before her arrival, or at least she has reason not to be glad she had failed to study Italian. The fact that the agent had most reason to do something in the past explains why she should have preferred it *then* as well as why she should *now* prefer to have done it in the past. In light of these observations we may formulate the following underlying principle:

Any Time Preference (ATP):

If a person has most reason (to ϕ at t_1), then she has most reason [to prefer overall (to ϕ at t_1) at any t]

The principle says that our overall preferences with regard to our actions should align with what we have most reason to do at the time of action. Thus, if I have most reason to stay home on Monday night, then on Sunday night I have most reason to prefer overall that on Monday night I will stay home; on Monday night I have most reason to prefer overall to stay home; and on Tuesday night I have most reason to prefer overall that I had stayed home on Monday night. In other words, the fact that I have most reason to stay home on Monday night entails not only what I should prefer overall on Monday night, but also what I should prefer overall at *any other time* with regards to my action on Monday night. This is the Any Time Preference (ATP) principle.

If the mother should not have had a child at a young age, then it would follow from ATP that the mother should now prefer overall not to have had a child at a young age. If the parents should have given their deaf infant the implant that would have allowed him to hear, then, according to ATP, they should later prefer overall to have done so. ATP implies that agents should prefer overall not to have done actions they had most reason not to do.

ATP seems so plausible that we often take it for granted. But what is the basis of this principle? I believe ATP gains its plausibility from a further principle that may seem to entail it.

Preference at Time of Action (PTA):

If a person has most reason (to ϕ at t_1), then she has most reason [to prefer overall (to ϕ at t_1) at t_1].

PTA says that *when* a person has most reason to do something, she should prefer overall to do it. Of course, much of the plausibility of the principle turns on what is meant by “preference.” For one, we can arguably have some preference for A while having overall preference for B, where B is incompatible with A. The alarm clock rings and I have some (considerable) preference for staying in bed, but I believe I have most reason to get up and, accordingly, prefer *overall* to get up. PTA does not require that we prefer *only* the course of action we have most reason to pursue; it requires that we prefer it *overall*.

Still, it may seem quite reasonable to wish I did not have to wake up early to go to work. In such a case I prefer overall not to have most reason to wake up early. This overall preference is quite reasonable, but it is also compatible with PTA. We can prefer

not to have most reason to wake up early and still prefer to wake up early given that we *do* have most reason to do so.

Now suppose I have an overall preference to stay in bed but I also have most reason to wake up. This would be a violation of PTA, but why is it objectionable? This is where our particular understanding of preferences becomes crucial. For it is quite reasonable to be saddened by, or dissatisfied with, what one has most reason to do. We often do things that sadden or pain us because we believe we ought to do them. In this sense, it may be reasonable to have an overall preference not to do what one has most reason to do. I discuss this plausible notion of regret in the next section of this chapter, but I think something stronger is involved in the cases we are considering.

If we understand a person's overall preference as what he or she is effectively motivated to do, then by having an overall preference to stay in bed I am effectively motivated to stay in bed and therefore to act contrary to what I have most reason to do, which is to wake up. We might put the point by saying that when one has an overall preference not to ϕ one is already in the process of not ϕ -ing. But if not ϕ -ing is contrary to what one has most reason to do, then by having an overall preference not to ϕ one is already in the process of doing what one has most reason not to do.

And now we are a short way from a paradigmatic case of weakness of will, in which, believing I have most reason to wake up, I prefer to stay in bed and slam the snooze button. PTA explains why believing I have most reason to ϕ commits me to the overall preference to ϕ , where "overall preference" is tantamount to effective motivation. Thus, an agent's overall preference with regard to possible courses of action should

accord with the course of action the agent has most reason to pursue. If PTA entails ATP then we have an explanation of the plausibility of ATP.

But PTA does not entail ATP. It does not follow from the fact that an agent should prefer overall to ϕ at the moment in which she ought to ϕ that an agent should also prefer overall to have ϕ -ed after the moment she ought to have ϕ -ed, or that the agent should prefer overall that she will ϕ before the moment she ought to ϕ . Unlike a lack of effective motivation in the moment of action, a lack of effective motivation to do what one ought to do in the future or what one ought to have done in the past does not imply a failure to perform the required action. So we may ask: What kind of failure might we attribute to someone who wishes she did not do something she ought to have done, or to someone who is glad she did something she ought not to have done?

It would not help here to say that if a person does not regret an action she should not have done then she fails to recognize that she should not have done the action. For this inference is precisely what we need to explain: why should we infer from the fact that a rational person overall prefers an action that she believes it was justified? What is the source of the rational requirement to prefer to have done what we believe we ought to have done?

Thomas Nagel has addressed this question in *The Possibility of Altruism*. According to Nagel, judgments about practical reason have “motivational content.” That is to say, when an agent judges that an action is justified, this is “sufficient, in the absence of contrary influences, to explain the appropriate action, or the desire or willingness to perform it” (Nagel 1970, 67.) PTA is seen here as a requirement of rational agency: the

fact that an agent judges that she ought to ϕ , together with the fact that a rational agent complies with PTA, explains why, insofar as the agent is rational, she prefers overall to ϕ , that is, she is motivated to ϕ .

Nagel then argues that seeing oneself as one, equally real person over time involves being motivated to do what one has reason to do, whether the relevant time of action is the future, the present, or the past. Thus, the motivational content of judgments about practical reason has different implications depending on the agent's temporal relation to the action in question. When one regards an undertaking as justified, "this is sufficient to explain one's wanting it to happen, be happening, or have happened" (Nagel 1970, 70).

Thus, Nagel's view suggests that we get from PTA to ATP via an assumption about the motivational implications of seeing an action as one's own. A person who overall prefers to have done something she believes she had most reason *not* to do is alienated from her past self: she does not recognize her past action as her own. To recognize one's past action as one's own is to be motivated by what one believes one had reason to do. Regret is therefore explained by the motivational content of the agent's judgment that he or she had in the past most reason to perform a different action than the action he or she actually performed. To recognize her mistaken decision as her own, the young mother must prefer not to have made it, and in this sense regret it. Is this a convincing explanation of ATP?

Nagel's proposal turns on the appeal of the idea that taking PTA to entail ATP is a necessary condition for seeing oneself as one, equally real person over time. John Rawls,

for one, has found this idea attractive and elaborated on what it involves. In the third part of *Theory of Justice*, Rawls discusses the question of goodness for an individual and the rational basis of regret. Rawls adopts Josiah Royce's notion of a *plan of life*, according to which "an individual says who he is by describing his purposes and causes, what he intends to do in his life" (Rawls 1971/1993, 408/358). A person, then, "may be regarded as a human life lived according to a plan" (ibid.) and a person's identity is given by his or her specific plan of life. The question of goodness for an individual, according to Rawls, is the question of what is a rational plan of life? That is, what plan of life should a person adopt? Rawls then argues that seeing oneself as one, enduring individual over time involves following one's rational life plan, or, at least, the life plan that is most likely to be one's rational life plan. If a person realizes that she has not been following the best life plan for her, or the life plan she should have pursued, then she has reason to regret her choice of life plan, even if she should not reproach herself for her choice, which might have been subjectively justified given the information available to her (Rawls 1971/1999, 422/370).

On this view, to see oneself as one, equally real person over time is to see all of one's choices—past, present, and future—as part of a single life plan that one ought to endorse. Thus, with every choice a rational person makes she endorses her plan of life, and since a person's plan of life includes her past and future choices, with every choice a person makes she endorses the past and future choices that her life plan prescribes. The notion of a rational plan of life connects PTA to ATP and explains Nagel's idea of seeing oneself as one, equally real person over time.

Rawls's suggestion is clearest in the case of small-scale plans. Normally, when I reserve a table at a restaurant I endorse a plan to have a good meal there. If I find that the food is not to my liking then I have reason to regret my choice of restaurant because it did not serve my present plan to have a good meal. If, despite the fact that I did not enjoy the meal, I affirm my choice of restaurant, say, because I am glad to have tried it, then I disavow my plan to have a good meal. The plan I pursue makes sense of the actions and choices that compose it. My attitudes toward my past and future actions and choices are determined by the fact that I endorse the plan as a whole. Rawls expands this idea to one's whole life. A failure to regret an action I should not have taken is a failure to prefer, in the present, the life plan that I should endorse. This violates PTA, which requires that I prefer in the moment of choice the option that I have most reason to choose, which is, in this case, my rational life plan.

But is it plausible to identify a person with his or her life plan? To understand who you are, is it sufficient to understand your current aims and causes? Are not our past failures and the persons we lost just as crucial for understanding who we are as our current projects and the persons we love? It seems to me that even the few persons who actually have a whole life plan worked out rarely pursue it without change. One, equally real person over time normally has different life plans at different moments in time. In this regard, Stuart Hampshire has offered the apt, following observation:

[A person] explains himself to himself by his history, but by the history as accompanied by unrealized possibilities on both sides of the track of actual events. His individual nature, and the quality of his life, do not depend only on the bare log-book of events and actions. His character and the quality of his experience emerge in the possibilities that were real possibilities for him, which he considered and rejected for some reason or other. (Hampshire 1989, 101)

With the same concern in mind, Bernard Williams has argued that the model of a rational life plan “is that of one’s life as a rectangle, so to speak, presented all at once and to be optimally filled in” (Williams 1981b, 33). The model, says Williams,

implicitly ignores the obvious fact that what one does and the sort of life one leads condition one’s later desires and judgments ... So there is no set of preferences both fixed and relevant, relative to which the various fillings of my life-space can be compared. (Williams 1981b, 34)

The point of Williams’s criticism, as I understand it, is that a person’s choice of life plan reflects the person’s deepest attachments and concerns at a given moment in time. We plan our lives in light of the relationships, values, projects, and activities that give meaning to our lives and define who we are as a person. Even if, in light of a person’s attachments and concerns, he or she should choose plan B rather than plan A, there is no sense in which the person should choose plan B independently of his or her attachments and concerns. A rational life plan is a plan for an actual person, given who he or she actually is, not a plan for the person he or she should have been. And since a person’s attachments and concerns normally change over the course of his or her life, a person’s rational life plan normally changes as well, despite the fact that the person is an equally real person over time. So when a change occurs in our attachments, it may change our perspective on our past choices and actions even if it does not change what we had most reason to do at the time.

This suggests that, contrary to Nagel’s and Rawls’s claim, a person may see him- or herself as one, equally real person over time even if, given the person’s present attachments, she affirms actions she had most reason not to undertake in the past. So even if every particular choice entails the endorsement of a plan that goes backward as well as forward in time and commits one to preferring certain past and future choices, this plan is

determined by one's present attachments and concerns, which may and normally do change with time.

Williams has concluded from this that whether an action is justified may be determined by subsequent events, which determine the agent's later attachments (Williams 1981b, 24). He explains:

The perspective of deliberative choice on one's life is constitutively *from here*. Correspondingly the perspective of assessment with greater knowledge is necessarily *from there*, and not only can I not guarantee how factually it will then be, but I cannot ultimately guarantee from what standpoint of assessment my major and most fundamental regrets will be. (Williams 1981b, 35)

But even if Williams is right, as I think he is, about the relevance of a person's subsequent attachments to her regrets, it does not follow that the person's subsequent attachments determine whether the past action was justified. To be sure, in hindsight we may acquire relevant information and understanding that we did not have at the time of choice and thereby come to a different assessment of what we had most reason to do. This is the point I made earlier about the difference between subjective and objective justification. But Williams's idea is that later events might *make* the past choice justified, not merely *reveal* that it was. If the young girl waits and has a child only when she is twenty-four years-old, she may come to reasonably affirm her decision to wait; if she has a child at fourteen she may come to reasonably affirm that choice as well. Williams's suggestion implies that the girl's later affirmation renders her choice justified *retroactively*, not that the girl's affirmation is an expression of her judgment that the choice was justified to begin with. According to Williams, what we had most reason to do in the past tracks what we presently prefer to have done.

This is a highly perplexing suggestion. For it seems to imply that the justificatory status of the girl's choice will be determined by her choice. In asking herself, at fourteen, whether she should have a child, it might seem that the girl should aim to satisfy the preferences of her future self. But since the preferences of her future self will be determined *by* what she does at fourteen, the question of what she should do at fourteen is downstream from the question of what she *will* do at fourteen. Her question, "What ought I choose?" cannot arise before she chooses. It is an essentially retroactive question. But of course the question *can* arise in the moment of choice. Indeed, that is when the question is most pressing, and that is why it is essential for rational agency. The idea that some actions, or maybe *all* actions, are justified retroactively, upends practical deliberation, which is essentially prospective.²¹

Moreover, retroactive justification seems to put too much weight on the agent's later point of view. While hindsight might allow for a better understanding of the relevant facts—some of which were not yet determined at the time the choice was made—an agent's later preferences, attachments, and concerns need not have a greater say about what should be done than the agent's preferences, attachments, and concerns at the time of action. Insofar as there is no difference in available information and understanding—that is, insofar as there is no epistemic difference between the points of view—there is no reason to grant our future self the final word about what our present self should do.

Fortunately, Williams's insight about the crucial role of attachments in determining one's reasons for regret leads to a different, more plausible conclusion. If we

²¹ Compare Wallace 2013, 163-168.

accept the claim that what a person has most reason to do is determined by the person's present attachments, then the person's later attachments may give her most reason to affirm what she had most reason not to do in the past. Thus, whether or not a choice is justified, a person may reasonably regret or affirm it on the basis of his or her later attachments. I therefore propose the following alternative to ATP:

Preference by Attachment (PBA):

If a person has most reason (to ϕ at t_1), then she has most reason to prefer [(to ϕ at t_1) at any t in which she has the same relevant attachments as the ones she has at t_1].

PBA explains regret in terms of the agent's attachments. If I have most reason to stay home on Monday night and write rather than go out to a movie, then for as long as I have the same relevant attachments I had on Monday night I will have reason to regret going out to a movie. But if my attachments change, then I might not have reason to regret going out to a movie. Suppose I had most reason to stay home to work on my stamp collection but went out to a movie. However, I have since sold my collection, and do not collect stamps anymore. PBA allows for the plausible judgment that I now lack reason to regret going out to a movie. The change in my attachments accounts for the fact that I may reasonably prefer to have made my past mistake.

This is also what happens in the case of the young mother and the parents of a deaf infant. At fourteen, the girl had most reason not to have a child given her attachments, but once she has a child, her attachments may change in a way that is relevant to her attitude toward the choice she had made. The requirement to regret her mistaken choice is nullified due to her new attachment to her child. Similarly, the parents

have most reason to give their infant an implant that would allow him to hear. But if they do not do so, and their child grows up to be a happy, fulfilled deaf adult, then the parents' attachment to their child as he actually is nullifies the rational requirement to regret their mistaken choice, which presupposed different attachments.

We now have an explanation of the difference between the easy cases, where one's mistakes should be regretted, and the difficult cases, where one may affirm past mistakes. Wallace is right to emphasize that the young girl has new reasons once she becomes a mother, reasons she did not have beforehand. But what explains the fact that the young mother is not required to regret her mistake is not that she has new reasons that she did not have before, nor that the action was justified by later events, as Williams argues, but that *regret about past choices and actions is rationally determined by one's present attachments*. The drastic change in the young girl's attachments explains why she had most reason not to have a child at fourteen and yet need not regret her decision once she has a child whom she loves.

Our solution is in line with Wallace's own solution, which also emphasizes attachments. However, now we have an explanation of why the solution is not ad-hoc. For we have rejected the assumption that made it seem ad-hoc in the first place, namely, the assumption that our overall preferences with regard to our past choices should accord with what we had most reason to choose. Instead, I have argued, our overall preferences with regard to our past choices should accord with what we had most reason to choose *insofar* as the attachments relevant to our choices have not changed significantly.

PBA also accounts for our preferences with regard to the future. The fact that the young mother has most reason to be glad she decided to have a child does not imply that she should have preferred to have a child to begin with. What she had most reason to prefer was determined by the attachments she had then, not by the attachments she formed later in life. Indeed, the reasons she has as a young parent account for why she had most reason to postpone motherhood. Or consider another example. In my youth I played the upright bass and studied music. I spent most of my days listening, playing, arranging, and composing music. Later I decided to quit music in order to study philosophy. Even if I made the right choice in quitting music, this does not imply that all along I had reason to prefer that I will quit music. When I played and studied music I reasonably preferred that I would continue to do so: valuing music meant wishing to continue to be enthralled by it. More generally, the fact that I will have most reason to do something does not imply that I should currently prefer that I will do it, for the attachment that will give me the reason to do it might not be in place yet. Thus, our present preferences are hostage neither to our past attachments nor to our future ones.

Someone might worry that the claim that a person's attachments determine what he or she has reason to do commits us to blatant relativism. What if the person is attached to his job as a hit man for the mafia?²² Does such person have most reason to be a good murderer? My distinction between reasons *for* attachment and reasons *of* attachment may help us address this worry.²³ What a person has reason to do is determined by the

²² See Korsgaard 1996, 254-255.

²³ In fact, in the previous chapter I draw the distinction in terms of reasons *for* valuing and reasons *of* valuing. But I also argue there that valuing an object and being attached to it may come to one and the same thing. In the text I use "attachment" in keeping with the terminology used in the literature on regret, e.g., in Wallace 2013 and Harman 2009.

person's reasons *of* attachment, but, as I argued in the previous chapter, reasons *for* attachment are enabling conditions for reasons *of* attachment. In order to have reasons *of* attachment with regard to a person or thing, the person or thing must be worthy of attachment. Philosophers often distinguish between *motivating reasons* and *normative reasons*, where motivating reasons may explain agents' actions even if they lack justificatory force. For instance, I might, out of envy, take the fact that doing X would undermine a colleague's career as a reason to do it. This would be my motivating reason, but it is not a good reason, so undermining my colleague's career would not be a normative reason for doing X. I propose we draw a similar distinction between *normative attachments* and *motivating attachments*. A person may be motivationally attached to his job as a hit man, but that would not be a normative attachment, which gives him normative reasons *of* attachment. If an agent takes him- or herself to have reasons *of* attachment with regard to a person or thing that she ought not be attached to, then, despite having a motivating attachment, the person lacks a normative attachment. We may therefore clarify PBA by saying that what a person has most reason to do is determined by the person's normative attachments, and that whether a person should prefer to have done an action, or to do it in the future, depends on the person's present normative (rather than motivating) attachments.

To summarize, PBA takes from Nagel the idea that we do not see our choices and actions in isolation but in terms of who we are as persons; it takes from Rawls the idea that we understand who we are as persons in terms of our goals and causes; but it rejects the idea that we can only recognize an action or choice as our own if we see it as part of our current life plan. A person's *current* (normative) attachments determine his or her

rational plan and make sense of the way the person reacts to his or her past and future choices and actions. There is no life plan that a person ought to adopt independently of her (normative) attachments. And although regret looks at the past, it always does so from the present, never from the point of view of eternity, so to speak. By replacing ATP with PBA we can explain why rational agents may sometimes reasonably affirm past mistakes that they acknowledge as such.

But there is a lingering doubt with regard to the young mother and the parents of the deaf child. It seems that the agents' mistaken decisions in the past are still relevant to their present lives. The young mother would have been a better, more able and mature parent had she waited to have a child at a later stage in life. She would have also had more opportunities in life to explore her talents and interests. The parents would have provided their child with a better life had they given him the implant that would have enabled him to hear. Should not these facts resonate in the agents' present attitudes toward their past? Can they reasonably and wholeheartedly affirm their past mistakes in light of what could have been? Or consider the matter more generally. We said that one's present attachments do not define who one is, for, as Hampshire put it, to understand a person we must understand his or her "unrealized possibilities on both sides of the track of actual events." The life plans we once had are still relevant to the persons we are today. But how do these lost possibilities express themselves in our actual, realized lives? The discussion in the following section will enable us to answer these questions.

3. Regretting a Loss

Thus far, I have proposed that there are reasons for regret, which are given by the object of regret, and that in light of a change in our attachments we might reasonably affirm an action we had most reason not to undertake. In this section I argue that reasonable regret is often not a response to mistakes at all, but rather a response to loss. Choices that involve loss may warrant regret even when the choice itself was justified or required. For, very often, what underlies regret is a not a longing for the decision we did not make, but a longing for the path we did not pursue.

*

To bring into view the possibility of reasonable regret about justified choices, consider optional choice situations, in which an agent has sufficient reason to choose either of several options. Optional choice situations can involve low or high stakes. In choosing what to do tonight someone might have several worthy options; but someone might have several worthy career options, too. The higher the stakes are, the more certain we want to be that we did not miss any relevant fact that might tilt the scale in favor of one of the options. But, presumably, we can and do face genuinely optional, high-stakes choice situations.²⁴

²⁴ Scanlon (2014) and Raz (1999a) argue for the prevalence of optional choices, where reason renders various alternatives “eligible,” to use Raz’s parlance. Widespread optionality is part of what Raz calls *the classical conception of agency*, according to which in the paradigmatic case of action reason renders options eligible. Raz contrasts this conception of agency with *the rationalist conception of agency*, according to which in the paradigmatic case of action reason *requires* the selection of one of the options.

Take the example of choosing between graduate schools. Perhaps I have sufficient reason to choose either of two schools. I consider the various advantages and disadvantages of each school and come to the conclusion that no matter which school I choose, my choice would be rationally justified because each option is supported by sufficient reason and neither by conclusive reason. Still, there may be significant differences between the options. Each school may lack the qualities and advantages of the other. Furthermore, each choice might lead me down a different life path—I will form different friendships, different teachers will influence my thinking in different ways, etc. It seems to me that whichever choice I make, I might later have reason to prefer overall that I chose otherwise, not because the choice I made was unjustified, but because I had good and sufficient reason to choose differently. A similar verdict is true of choices with lower stakes. I might have sufficient reason either to stay home and read a book or to go out to a movie. Still, the fact that the options are worthy for different reasons may later give me reason to wish overall that I did not choose the option I in fact chose. If this is true, then we may reasonably regret justified choices.

It is tempting to resist this line of reasoning and insist that it would not be reasonable to regret a rationally justified choice, a choice one had most reason to make. Even if there are some ways in which the option I did not choose would have been better or more worthy than the one I did choose, it might be argued that since I had sufficient reason to make the choice I actually made, I lack reason to have an overall preference for the alternative. If both options were justified then the agent lacks reason to have an overall preference in favor of one of the two options and therefore lacks reason to retroactively prefer one of the two options. In this way, we might wish to defend the idea

that we should not regret optional choices and that regret is reasonable only as a response to past mistakes.

But I do not think it is plausible to hold that we have reason to prefer an alternative only when we have conclusive reason to choose it over all other options. To see why I find such contention implausible, consider the distinction between *picking* and *choosing*. Sidney Morgenbesser and Edna Ullmann-Margalit have distinguished between “selection situations,” as they call them, where there is and selection situations where there is not a relevant difference between the alternatives. They call the former kind of selection situations *choosing situations* and the latter kind *picking situations* (Morgenbesser and Ullmann-Margalit 1977).

Morgenbesser and Ullmann-Margalit understand the distinction between picking and choosing in terms of the agent’s preferences. They hold that picking situations are ones in which the agent is indifferent between the options and choosing situations are ones in which the agent prefers one option to the others. But Morgenbesser and Ullmann-Margalit assume that what accounts for the agent’s preferences are facts about the options. We might therefore say that picking situations are situations where the facts about the alternatives give the agent no reason at all to prefer one option to the others, whereas choosing situations are ones in which the facts about the alternatives do give the agent reason to prefer one option to the others. Buridan’s ass faces a picking situation, for it has no reason at all to prefer one stack of hay to another. Were one stack of hay closer or larger than the other, the ass would face a choosing situation, in which there is a

relevant difference, i.e., there is some fact (or facts) that counts in favor of one of the options and against the other.

But now we can see that there are two kinds of choosing situations. There may be *strong choosing situations*, in which facts about the alternatives give the agent conclusive reason to prefer one option. This is the case, for example, when one faces a choice between different means to the same end and the only relevant difference between the means is that one is more effective than the other. But there are also *weak choosing situations*, in which there are significant differences between the options and yet each option is supported by sufficient reason. Such is the case described above of choosing between graduate schools, or the case of choosing between going out to a movie and staying home to read. There is substantive difference between the facts that count in favor of going out and the facts that count in favor of staying home, but one would be justified in choosing either option.

In picking situations, there is no substantive reason to prefer one option to the other. Rather, selection is made on the basis of higher-order reasons to prefer *some* option.²⁵ But in weak choosing situations, lack of indifference is reasonable: there is substantive difference between the rationally permissible options. In other words, when I have sufficient reason to go out and sufficient reason to stay home it would make sense to have a persistent preference in favor of one of the options. This is not the case when I choose between two cans of Campbell's Tomato Soup. Unlike picking situations, weak

²⁵ For accounts of higher-order reasons for selection in picking situations, see Scanlon 2004 and Kolodny 2011.

choosing situations allow for reasonable discrimination between justified, choiceworthy, rationally permissible options.

If there are weak choosing situations, as seems very plausible, then it does not follow from the fact that one has sufficient reason to choose either option that one should be indifferent between the options.²⁶ Rather, in weak choosing situations the agent may reasonably prefer overall either option to the other. Thus, in retrospect, too, it would be reasonable to regret—to have an overall preference against—a justified choice if there were substantive reasons to favor the alternative. Furthermore, in line with our conclusion from the previous section, we may say that whether we presently have reason to regret a justified choice is determined by our current normative attachments, which may be relevantly different from the attachments we had at the time of choice.

It therefore seems that, at least in weak choosing situations, we may reasonably regret an action we had sufficient reason to undertake. The explanation for this may be found in the difference between picking situations and weak choosing situations. In picking situations the loss of the option not chosen is *insignificant*; in weak choosing situations the loss of the option not chosen is *significant*. The significance of the loss is reflected in the different types of reasons that count in favor of each option. When choosing between different cans of Campbell Tomato Soup, the same type of reasons support each option. But this is not normally the case when we choose, for example, whether to stay home or go out. We may therefore say that weak choosing situations involve loss, while picking situations do not. So it seems that what makes sense of a

²⁶ For a view according to which any choosing situation is a strong choosing situation, see Regan 1997.

person's preference that she did not do something she had sufficient reason to do is the loss of the option that has not been pursued. In weak choosing situations we do not regret our mistakes; we regret our justified losses.

Once we focus on the idea of loss, we may notice that even strong choosing situations, in which the agent has conclusive reason to choose a certain option, may involve loss. Suppose I have conclusive reason to stay home and write rather than go out tonight. Still, there are reasons in favor of going out, and those reasons are different in type from the reasons that count in favor of staying home. Different types of reasons might count in favor of each option even if only one option is supported by conclusive reason. Other times, strong choosing situations might involve no loss, as in the case where I choose the most effective means over the less effective means or in the case where one can of Campbell's Tomato Soup is more expensive than another and all else is equal. The same types of reasons count in favor of each option, but only one option is supported by conclusive reason. The question of whether a choice involves loss is orthogonal to the question of whether the choice is supported by conclusive or merely sufficient reason.

We may therefore introduce a further distinction, a distinction between *cost situations* and *loss situations*. The idea is that when a choice between different options does not involve significant loss, it might still involve a cost. When two options are supported by the same kind of reasons, and are therefore worthy in the same way, but only one option is supported by conclusive reason, then choosing the option that we have conclusive reason *not* to choose, would involve a cost, but no loss. For we could have had

more of the same thing, as it were. *Cost situations* present the agent with options supported by reasons of the same type, while *loss situations* present the agent with options supported by reasons of different types. Picking situations are always *cost situations*: indifference is called for due to the fact that either option is justified and both are justified by the same type of considerations. Weak choosing situations are always *loss situations*: they always involve several justified options supported by reasons of different types. Finally, strong choosing situations are either *cost situations* or *loss situations*, depending on whether the reasons that support one option are different in type from the reasons that support another.²⁷

Now, a baffling possibility arises. I have argued that loss accounts for regret in weak choosing situations, and that strong choosing situations may also involve significant, inevitable loss. It therefore seems that even choices made in strong choosing situations and supported by conclusive reason may be reasonably regretted. For even a choice we have conclusive reason to make may involve loss. But unlike weak choosing situations, in strong choosing situations regretting a justified choice seems to imply wishing to have done something one ought *not* have done. How might it be reasonable to prefer overall to have acted contrary to what one had most reason to do? How might it be reasonable to regret the only correct, reasonable, required choice one could have made?

This worry may be addressed by making a distinction in the object of regret. To say that a choice was justified, or that it was supported by conclusive reason, is to say that it was an appropriate choice given the alternatives. Therefore, when we regret an

²⁷ Monist views about value, which hold that there is, at bottom, only one substantive value, construe all choice situations as cost situations, where no significant loss is involved. It seems to me that the fact of loss counts against monism about value.

unjustified choice as such, we regret it in light of the alternatives and prefer to have chosen a different alternative. However, even alternatives that are justifiably discarded may involve significant loss. Therefore, when we regret *the loss of an option* we might not regret *the choice to forgo it*, for in the circumstance the choice might have been justified or even required. What we regret is the absence of the option we have not pursued. We wish that we had pursued the lost alternative, not that we had pursued it instead of the alternative we actually pursued. We regret the loss, not the decision that brought it about.

The word “regret” is somewhat misleading here because it may naturally suggest that an agent believes she should have taken a different path than the one she actually took. Talk of “preferences” is also misleading in this way. To say that one prefers not to have done what one did suggests that one prefers to have done something else instead. But once we peel-off the word “regret,” avoid “preferences,” and stick to “wishes,” we can distinguish between [wishing I ϕ -ed instead of ξ -ed] and [wishing I ϕ -ed]. Wishing that some course of events were actual does not commit the agent to preferring it to an alternative; it only commits the agent to preferring its actuality to its absence.

Consider again an example from the previous section. In my youth I played the upright bass and studied music. At a certain point, I decided to quit music in order to study philosophy. Today, many years later, I sometimes think back on that decision. When I visit my parents’ house and see my old bass standing in the corner of the study, or when I go to a concert and listen to the musicians play some tune I once played myself, I feel a sense of loss that reflects the life I could have pursued but never did. Even

in moments like these I do not regret my decision to study philosophy. For various reasons, it was the right choice to make. Still, I wish that I continued to play and compose music. The focus of my regret is not my decision to forgo a life of music; it is the life of music that I forwent. What I wish for is music, not the choice to have pursued it.

Take a different, more striking example. Peter Baumann and Monika Betzler discuss Winston Churchill's alleged decision to let the Germans bomb Coventry to prevent much greater harm (Baumann and Betzler 2004, 8). The story has been disputed so let me briefly give some background. On the night of November 14, 1940, five hundred and sixty-eight residents of Coventry were killed, eight hundred and sixty-three were severely injured, and three hundred and ninety-three sustained lesser injuries. F. W. Winterbotham has claimed, in a book published in 1974, that the British government had advance warning of the attack on Coventry thanks to intercepted German radio transmissions (Winterbotham 1974). Winterbotham further claimed that Churchill decided to allow the attack in order not to give the Germans cause to suspect that the British have deciphered their code. Later historians have rejected both hypotheses and it is now generally agreed that while the British government had prior knowledge of a German air raid, the exact target of the raid was unknown.

Were Winterbotham's hypotheses true, and Churchill were to know of the raid on Coventry in advance and decide (justifiably, let us say) to allow it, we would expect such decision to give rise to agonizing regret. Indeed, it would arguably be objectionable of Churchill to be indifferent to the bombing of Coventry even if his decision were the right one. The point is not that Churchill should regret finding himself in these circumstances—although he should regret that, too—but that he should regret the fact

that he *actually* let the Germans bomb Coventry, that he should regret the lives that were lost and the grave harm brought about innocent civilians. Given that, in our hypothetical scenario, Churchill had to choose between two gruesome options, one of which worse than the other, letting the Germans bomb Coventry and kill many civilians was the right decision to make. And still, the deaths of many innocent men, women, and children are deeply regrettable if anything is. Churchill's regret would have expressed his appreciation of the loss involved in the justified choice he made.

Thus, even with regard to choices we had conclusive reason to make, we may reasonably regret the lost alternatives, even if we cannot regret the decision itself. In contrast, when faced with weak choosing situations, in which we have substantive reason to prefer overall either alternative, we might reasonably regret our decision as well as the lost alternative. And, finally, in picking situations we may reasonably regret neither the decision nor the lost alternative, for there is no significant loss involved in the justified choice.

Now we may address the lingering question from the previous section, with regard to the young mother and the parents. We said that these agents may reasonably affirm their past mistakes given their present attachments. But we can see now that despite their affirmative attitudes toward their past decisions, they may presently regret the lost options. The mother might regret not being a better parent, and not pursuing various opportunities she would have had had she waited to have a child at a later stage in life. The parents may regret their child's inability to enjoy sound, to hear their voices, and to participate in activities that require hearing. Such regret does not commit the agents to regretting their past decisions, but only expresses their appreciation of the loss involved

in those decisions. The young mother and the parents may reasonably affirm their decisions while reasonably regretting the loss they involved. In this way, even when they are not regrettable, past mistakes may remain relevant in the present.

However, not every significant loss is a loss we should regret. The life plans and attachments we left behind might fade and disappear into oblivion—places we used to frequent, hobbies we gave up, friends who moved away, projects that never came to fruition. Our lives involve more loss than we can fathom. The losses we do regret express our lingering attachments and moral commitments. This is why I still feel the loss involved in my choice to forgo music, and why the hypothetical Churchill should have continued to regret the loss of life in Coventry. Lost attachments may and often do accompany us in the present as latent, unrealized possibilities, manifested in our sense of loss and regret. The attachments we have failed to pursue, or had reason to give up, are our *unrealized attachments*, and, as Hampshire suggests, they are as much a part of the persons we are today as those attachments that determine the visible structure of our lives.

4. Regretting an Attachment

In the previous two sections we have seen that an agent may reasonably affirm a past mistake while regretting the loss it involved. I would now like to apply the lessons of the previous sections to cases of political concern that implicate us, via our attachments, in deeply regrettable actions and events. First, consider “the bourgeois predicament.” R. Jay Wallace argues that the kinds of projects, relationships, and activities that are sources of meaning in the lives of many of us, are noncontingently conditioned on past and

present wrongs of great magnitude. For example, Wallace argues that oppression and deprivation have played a crucial role in the history of American universities, and that the current level of academic research in these institutions would not have been possible in a more equitable, global state of affairs. Wallace believes that such facts lead to the bourgeois predicament, which

stems from the role of objectionable historical and material conditions in making possible the attachments that, as a matter of fact, fill out our lives and render them worthy of retrospective affirmation. The very things that give us reason to be glad that we have lived are conditioned by circumstances that we can only wish had been otherwise. (Wallace 2013, 227-228)

Wallace's description of the bourgeois predicament might also fit the case of Israelis' attitudes toward the Palestinian Nakba, that is, the displacement and dispossession of hundreds of thousands of Palestinians during the Israeli War of Independence. Israelis fear that acknowledging the disaster Israel had wreaked on the Palestinians would commit them to regretting the attachments that give meaning to their lives. And, in the same way, many Israelis are reluctant to acknowledge the injustice of the current Israeli military rule over Palestinians in the West Bank because they fear that it is necessary for the attachments that are most important to them. Like those in the bourgeois predicament, Israelis struggle to come to terms with the fact that their most meaningful attachments rely on grave, systemic wrongs. We may call this the "Israeli predicament."

Both the bourgeois predicament and the Israeli predicament have a historical aspect and a contemporary aspect. The historical aspect involves past wrongs that led to the valuable attachments that give meaning to the lives of those in the relevant predicament. The contemporary aspect involves the current wrongs that sustain these

attachments. The two aspects should be clearly distinguished, for they raise different problems and call for different responses.

Consider the historical aspect first. Let us accept, if only for the sake of argument, that grave wrongs toward Palestinians were a necessary part of the founding of the State of Israel, and that historical practices of social hierarchy and exclusion were necessary for the current level of research at top academic institutions in the US. Let us also accept that the consequences of those past wrongs include much that is valuable for its own sake and that gives meaning to the lives of many individuals today. The question is whether one should prefer, today, that the wrongs in question did not occur, and, if so, does such preference undermine one's attachments to the consequences of these wrongs?

Wallace argues that if the historical conditions in question were necessary for the existence of the attachments on the basis of which a person affirms his or her life, then the person is committed to affirming these historical conditions, despite the fact that they are deeply regrettable. This is because, according to Wallace, wishing that the historical conditions were otherwise, commits the person to preferring that the objects of attachment that these conditions made possible did not exist, and this, in turn, undermines the basis of the person's affirmative attitude toward his or her life. This is an instance of what Wallace calls "the affirmation dynamic," which takes us from the affirmation of the object of attachment to the affirmation of the necessary conditions for its existence (Wallace 2013, 222).

But why should regretting the objectionable historical conditions commit us to regretting the attachments that give meaning to our lives? Wallace understands regret in

terms of on-balance preferences with regard to the past, analogous to intentions to have acted in a certain way, or to counterfactual commitments to act in a certain way in certain circumstances (Wallace 2013, 55-57). The question that regret answers, according to Wallace, is this: “would we, knowing what we know now about how things have since played out, bring it about that things were otherwise in the respect that we are focusing on, if it were in our power to do so?” (Wallace 2013, 62). Thus, Wallace’s notion of regret leads him to conclude that to regret the objectionable historical conditions is to have an intention-like attitude to prevent them had that been possible. Such hypothetical action would have blocked the chain of events that led to the attachments on the basis of which the agent affirms his or her life. So to regret the historical conditions commits one to preferring a world without the projects, pursuits, relationships, etc., that give meaning to one’s life. Since the past wrongs are deeply regrettable, and since we should acknowledge them as such, Wallace concludes that we face a condition of “deep ambivalence” (Wallace 2013, 249). The affirmation dynamic commits us to affirming historical conditions that we cannot but admit are deeply regrettable and should not have existed. Thus, “[o]ur plight as humans might be that we are condemned, in virtue of our attachment to life, to affirming conditions in the world that we cannot possibly regard as worthy of this attitude” (Wallace 2013, 255).

One might worry that the condition of “deep ambivalence” that, according to Wallace, is our plight, is nothing but a condition of rational failure, in which our intention-like preferences with regard to the past fail to track our judgments. Given Wallace’s contention that regretting the historical wrongs undermines the attachments that give meaning and worth to our lives, and given the fact that the historical wrongs *are*

regrettable, should we not conclude that our lives are meaningless and worthless? To be sure, it might be true that the fact that we would like our lives to be meaningful *psychologically* commits us to affirming that which should not be affirmed, but why should we concede that our wish to see our lives as meaningful *rationally* commits us to affirming that which should not be affirmed? And even if we cannot possibly go on with our everyday lives while accepting this stark conclusion, might we not have moments of reflective clarity in which we realize that our lives are meaningless? Indeed, one might think that Wallace's book aims to do exactly that: force us to acknowledge that we lack sufficient reason to affirm our lives as meaningful.

Wallace's dark conclusion is far-reaching. Wallace thinks that affirming our lives may commit us to affirming the entirety of world history, and suggests that this is the truth captured by Nietzsche's doctrine of the eternal recurrence:

Only if we are prepared to will the totality of world history can we honestly adopt an attitude of unconditional affirmation toward our lives and the other things to which we are attached. Because for all we know, that attitude already commits us to affirming as well the most catastrophic and egregious aspects of the larger histories in which our lives are caught up. (Wallace 2013, 253)

This claim is reminiscent of John Rawls's comment at the end of the introduction to *Political Liberalism*, where he joins Kant in holding that if justice is impossible it is hardly worthwhile for human beings to live on the earth (Rawls 1995, lx). This is why, Rawls insists, we *must assume* that justice is possible and consider what the possibility of justice entails. In a similar vein, Wallace's view might seem to propose that in light of humanity's morally wretched history we should regret that human beings have ever existed, even if we *must assume* that our lives are worthy of affirmation. Rawls asserts hope for the future; Wallace offers delusion about the past.

Resistant to Wallace's Nietzschean idea, Thomas Nagel suggests that

[i]n response to this extravagance it does not seem too defensive to suggest that our affirmation of anything, our own existence included, is bounded by a statute of limitations on its reach into the past. We can take much about the world that we have not created, good and bad, as simply given, and limit our affirmations and regrets to what is downstream from that. (Nagel 2014, 26)

In effect, Nagel suggests that Wallace's sweeping conclusion should be taken as evidence that his arguments are flawed. We should not concede that the life of most human beings are meaningless or worthless because they were made possible by grave wrongs in the past. Rather, we should doubt the apparent rational commitment that leads us from regretting the past wrongs that were necessary for the attachments that give meaning to our lives, to regretting those attachments and the lives they make meaningful. But having made this useful suggestion, Nagel does not go on to locate the flaw in Wallace's affirmation dynamic.

To find the flaw in the argument, let us look at why Wallace thinks we are committed to affirming our attachments in such a robust way. According to Wallace, for one to affirm one's life is "to prefer on balance that the series of events that constitutes the life should have occurred, as against the alternative that one should not have lived at all" (Wallace 2013, 148). Furthermore, a life is meaningful, when "the meanings that the life has achieved are substantial enough to make it worthy of affirmation on the whole, despite the presence of its many regrettable aspects" (Wallace 2013, 197). Thus, in order to affirm our lives in this unconditional way we must unconditionally affirm the objects of the attachments that give meaning to our lives. And to affirm the objects of our attachments unconditionally is "to prefer on balance that they are in fact part of the history of the world, taking into account the totality of things that they involved"

(Wallace 2013, 75). “The totality of things that they involved” refers here to the historical, normative, and constitutive necessary conditions for their existence. Wallace concludes that if the attachments that give meaning to our lives cannot be affirmed in this unconditional manner, then the attachments cannot be the basis of our affirmative attitude toward our lives, and we are rationally committed to preferring on balance that we did not live at all. And since the necessary conditions for the existence of the objects of our attachments include egregious wrongs, we cannot prefer on balance that these objects and the totality of things that they involved are part of the history of the world.

Wallace’s argument begins to unravel when we consider more carefully his notion of unconditional affirmation of objects of attachment. Why should we accept this notion of affirmation? Or, better yet, what would such affirmation even mean? In preferring that an object and all that it involves is part of world history, are we to prefer the actual world history, in which the object and all its historical necessary conditions exist, to *every* possible alternative history in which the object and its necessary conditions are absent? This interpretation of Wallace’s notion of unconditional affirmation would surely be too strong. For any given object in the actual world, we can imagine some possible world history that is less regrettable on balance than the actual world but from which the object in question is absent. Affirmation would be impossible. Perhaps, then, to affirm an object unconditionally only involves preferring the actual history, in which it and all its necessary conditions exist, to *some* possible world history where the object and its historical antecedents do not exist. But this interpretation is clearly too weak. For any given object we can always imagine alternative courses of history that contain more regrettable actions and events but in which the object and its historical antecedents do not

exist. The history in which the object exists cannot possibly be worse than *every* alternative history from which the object is absent. Affirmation would be trivial.

Perhaps, then, unconditional affirmation of an object implies preferring the actual world, in which the object and all that it necessarily involves exist, to the closest possible world in which the object and all that it necessarily involves do not exist. In other words, in order to imagine the alternative to the existence of the object of attachment, we should hold constant the actual world history and extract from it the object of attachment and all that was necessary for its existence. We should then compare the history of the world with the object to the history of the world without it and assess which one is on balance more regrettable.

Such comparison, it seems to me, is rather difficult to carry out. World War II was a deeply regrettable event, to say the least, and the world would have been a very different place if it did not occur. But it is difficult to say what the world would have been like without World War II, and, in particular, whether it would have been less regrettable than it actually is. The further back into history we look in search of necessary conditions for the existence of our attachments, the more difficult it is to grasp how history would have ensued in the absence of these past events. In fact, the most we can hope for is to consider the *likeliest* alternative history. But calculating likelihoods of hypothetical courses of history is a notoriously daunting, if not impossible task. How can we say whether the existence of the objects of our attachments and all that led to their existence is preferable to what would have been without them?

It seems to me that Wallace's notion of affirmation relies on what might be called a barrel-of-apples conception of world history. The history of the world is imagined as a barrel that contains many apples—persons, objects, actions, and events. When we judge that a particular apple is rotten, we prefer to throw it out of the barrel: one less rotten apple in the world! If the rotten apple happens to give meaning to our lives, then we are in a bind, because we have to throw ourselves out with it. Our desire to feel that we belong in the barrel/world blinds us to the rottenness of the apples that give meaning to our lives.

What is misleading about this picture is that when we take an apple out of a barrel we have one less apple in the barrel, but when we take an event out of world history we do not have one less event in world history, we have a different one. According to Wallace's affirmation dynamic, regretting that something exists involves preferring that the history of the world did not include it *as well as preferring the history that would have ensued in its absence*. To prefer that an apple were not in the barrel, we must also prefer its would-be replacement. That is, we must prefer the apple that would have been in the barrel if this rotten apple were not. Making such judgments seems not only difficult, but also irrelevant. We may deeply regret WWII without preferring the most likely alternative history.

Consider for a moment prospective preferences. Preferring one course of action to another does not require preferring the totality of future consequences of the first option to the totality of future consequences of the second. We normally have no clear grasp of the totality of future consequences of the options between which we choose, and yet we

may reasonably prefer one option to the other due to their *likely* consequences. In fact, preferring one option to another need not even involve preferring the likely consequences of one to those of the other. We may prefer an option due to obligations or commitments we have, or we may prefer it for its own sake and regardless or despite its likely consequences. If we reject the idea that the proper basis of our prospective preferences is an on balance weighing of regrettable consequence versus meaningful or valuable consequences, why think that our preferences with regard to the past should be couched in an on balance weighing of regrettable antecedents versus meaningful or valuable antecedents?

Wallace's mistake, I think, is that he runs together the two kinds of regret distinguished in the previous section: regretting that (X occurred rather than Y,) and regretting that (X occurred.) We can regret a past event, in the sense that we prefer that it did not occur, without being committed to preferring any alternative course of history. This is why we can regret WWII without plunging into frivolous historical speculation, and this is also why we can affirm the objects of our attachments while regretting their historical antecedents. Affirming the existence of American universities does not commit us to preferring their history of discrimination to a possible history in which they did not discriminate, and affirming the State of Israel does not commit me to preferring the Nakba to an alternative course of history in which the Nakba did not take place. The affirmation dynamic does not go through because it wrongly assumes that regret is always comparative and then forces us to choose between the objects of our attachments and the historical conditions that made them possible. But this is a false choice.

Nevertheless, our attachments do often implicate us in past regrettable events. The fact that the Nakba occurred is clearly relevant to my affirmation of the State of Israel. Similarly, some claim that without WWII the founding of the State of Israel would not have been possible. These facts—assuming the latter is a fact—are relevant to anyone who values Israel today. But I do not think that the relevance of these facts is best understood as grounding a commitment to regret the founding of Israel if one regrets WWII or the Nakba. Rather, the claims are relevant because a commitment to Israel and what it purportedly stands for requires that one appreciates the loss that made its founding possible.

More generally, the history of what we are attached to may be relevant to the content of the object of attachment. For example, learning that a person we love committed a great wrong before we met him or her may influence our understanding of this person and of the nature of our attachment to him or her. Our attachment might not withstand this discovery, since we may come to believe that we lack sufficient reason to be attached to this person in the way that we are. Naturally, our attachments to political or academic institutions involve greater portions of history. Attachments to universities that pride themselves on equal opportunity and to political institutions that pride themselves on being just may be undermined by discoveries about their discriminatory conduct in the past or about egregious crimes they committed in their wars of independence. The past is relevant to our affirmation of our attachments because it may indeed be part of what we are attached *to*. This suggests a distinction between the necessary historical conditions of the objects of our attachments, and the historical conditions of those objects that are part of the content of our attachment.

Thus, an American citizen may believe that the United States has a thoroughly glorious past, and this supposed fact might play a role in what this person sees as her reason for valuing the United States. When such a person discovers that American political institutions committed deplorable crimes against Native Americans, for example, the person is committed to reconsidering whether she has genuine reason to value the United States. Furthermore, even if the person does have reason for valuing the United States, the way in which she values her country must change, for such a person can no longer reasonably take pride in a crucial part of the American past.

The same is true in the case of attachment to American universities or to the State of Israel. Learning of past wrongs associated with these institutions might rationally force us to revoke our attachment to them. But there might still be much about these institutions that is worthy of attachment, and rather than revoke our attachments to them we may have reason to adjust our attitudes toward them so as to reflect their past misdoings. Past wrongs are relevant to our attachments because, and insofar as, they are part of what we are attached to, not because they were necessary for the coming about of the objects of our attachment.

This result leaves room for the possibility that discoveries about the past might undermine our present attachments. Wallace's worry is therefore warranted: we might face a conflict between our judgment of the past and our affirmative attitudes toward the objects of our attachments in the present. But even if we come to believe that our present attachments are not warranted, and that we should forgo them and forge new attachments, we need not conclude that our lives thus far have been meaningless or unworthy of

affirmation. There might still be much that was meaningful and valuable about these unjustified attachments, as misguided as they might have been. Our misconceptions about our lives do not drain them of worth.²⁸

The grain of truth in Wallace's worry is that our attachments normally require a kind of devotion that makes consideration of their appropriateness unavailable to us. A reluctance to face an ugly truth about the country we value, or about the institutions in which we partake, is a natural and even appropriate disposition given our devotion to them. Being devoted to something or someone often involves running the risk of being inappropriately devoted—that is, being devoted in an inappropriate way or to an object that does not warrant devotion. Such risk might very well be justified given the value of appropriate devotion. However, when the inappropriateness of our devotion becomes apparent to us, we cannot rationally ignore or deny it; we must confront its upshots, we must distance ourselves from the object of our attachment. Indeed, the mere acknowledgment of facts that count against our attachment to an object might already be a first step away from our devotion to it and towards a life that is made meaningful by different attachments. But acknowledgment of such facts might also be a first step on our way back to a more proper form of valuing. In this way, distancing oneself from what one values can be an act of devotion that enables the continuation of proper valuing, as in the case discussed above of the disenchanting American citizen.

²⁸ It seems to me that a view on which most human beings fail to meet a moral standard that is necessary in order to render their lives meaningful, is a view that itself fails a necessary moral standard: namely, it fails to acknowledge that, even in our morally wretched world, human life is meaningful. I do not mean to claim that Wallace subscribes to such a view, but his arguments may seem to suggest it.

In considering present wrongs, however, the conflict between devotion and the conditions for its appropriateness is more acute. Present systemic wrongs should be stopped in the present. Even if the military occupation of the West Bank provides better security to Israeli citizens and enables them to live a peaceful life—an assumption that is highly doubtful—the fact that it constitutes a perpetual attack on Palestinian society as a whole makes it wrong and condemnable. An Israeli citizen has conclusive reason to want Israel’s military rule to stop, and the Israeli government has conclusive reason to stop it. The reasons are essentially moral, but they are also connected to what Israelis purportedly value about Israel—its humanism and justice. However, if the end of Israel’s military rule implies that many of the valuable attachments that currently give meaning to the lives of Israelis will not be possible any more, how might Israelis coherently prefer it?

Or consider the current, global inequality that, presumably, accounts for the bourgeois predicament. Wallace says that many of the intrinsically valuable projects and relationships that this inequality makes possible would have to vanish in order to correct it. There is a conflict between our devotion to these projects and relationships and our moral obligation to end, or at least not sustain and take part in the condition of global inequality, which leaves a great segment of the world population in dire poverty. Assuming Wallace is correct, might we coherently prefer that global inequality will end rather than that our attachments continue?

The answer, it seems to me, is “yes.” If there is a serious moral objection to maintaining an attachment, then we have most reason to end it, even if it has thus far gave meaning to our lives. For example, if Wallace is right that destroying the lives of

whole nations is necessary in order to sustain contemporary philosophical research, then we should certainly be willing to end contemporary philosophical research. Our attachments are partly sustained by our recognition of our reasons *for* attachment in the relevant way. When we come to the conclusion that we have most reason *not* to be attached, this should put rational pressure on the attachment itself.

Similarly, the fact that contemporary Israeli life is sustained through a continuing assault on the Palestinians undermines the conditions for a genuinely normative attachment to all that contemporary Israeli life includes. This is why Israelis find it so difficult to accept the reality of Israel's military rule. To accept this reality would rationally commit them to drastically changing Israeli life and the conditions necessary for its existence so as to be worthy of attachment. It is precisely because Israelis recognize this implication that they adamantly refuse to acknowledge its antecedent. But the antecedent is, in fact, true, and the consequent does in fact follow.

Wallace seems to think that we have no choice but to affirm our attachments, since they are the basis of our attitude of affirmation toward our lives. But this seems to me to be another instance of a too narrow conception of the boundaries of the self. A life that is affirmed on the basis of one set of attachments might later be affirmed on the basis of a different set of attachments. Some attachments may simply run their course. Projects come to completion; goals are achieved. When they do, new attachments may take their place as a source of meaning in our lives. As I argued in the second section, to see oneself as one, equally real person over time is to recognize that one has survived and may still survive radical changes in the attachments that give meaning to one's life.

When our current attachments are conditioned on grave, systemic wrongs that continue to occur, I believe that we should not acquiesce in a condition of deep ambivalence, but rather strive to correct the wrongs in question while revising our current attachments or forming new, less objectionable ones. This is true of those in the bourgeois predicament as well as of those in the Israeli predicament. Some struggle to rescue their attachments from moral oblivion by changing their objectionable foundations, others forgo the attachments altogether in order to form new attachments that are not fraught with moral objections. But we should be careful not to condemn the attachments too quickly. Perhaps valuable academic work can be done in a world without deprivation, and perhaps Israeli culture and society may thrive without Israel's military rule in the Occupied Territories. Insofar as we are devoted to our attachments as well as to justice and morality, we should try to reconcile these aspects of our lives. But if such middle ground cannot be found, then, I believe, we should be willing to forgo our attachments and forge new ones, more compatible with morality and justice. We may still, of course, regret our lost attachments, and continue to wish we had not lost them; but we should lose these attachments, nonetheless.

Bibliography

Anderson, Elizabeth. (1993). *Value in Ethics and Economics*. Cambridge, Massachusetts: Harvard University Press.

Baron, Marcia. (1991). "Impartiality and Friendship." *Ethics* 101, 836-857.

Baumann, Peter; Betzler, Monika. (2004). "Introduction: Varieties of Practical Conflict and the Scope of Practical Reason." In *Practical Conflicts: New Philosophical Essays*, ed. Peter Baumann and Monika Betzler, 1-27. New York: Cambridge University Press.

Bittner, Rüdiger. (1992). "Is It Reasonable to Regret Things One Did?" *The Journal of Philosophy* 89, 262-273.

Brodsky, Joseph. (1986/2015). "On Mark Strand (1934-2014)." *New York Review of Books*, January 8, 2015.

Brown, Deborah. (1997). "The Right Method of Boy-Loving." In *Love Analyzed*, ed. Roger E. Lamb, 49-64. Boulder, Colorado: Westview Press.

Burnyeat, M. F. (1980). "Aristotle on Learning to Be Good." In *Essays on Aristotle's Ethics*, ed. Amélie Oksenberg Rorty, 69-92. Berkeley, California: University of California Press.

Cohen, G. A. (2013). "Rescuing Conservatism: A Defense of Existing Value (All Souls version)." In *Finding Oneself in the Other*, ed. Micheal Otsuka, 143-174. Princeton, New Jersey: Princeton University Press.

Dancy, Jonathan. (2004). *Ethics without Principles*. New York: Oxford University Press.

Darwall, Stephen. (2002). *Welfare and Rational Care*. Princeton, New Jersey: Princeton University Press.

- Emerson, Ralph Waldo. (1841/2010). "Self-Reliance." In *Self-Reliance and Other Essays*, 29-48. Nashville, Tennessee: American Renaissance.
- Frankfurt, Harry. (1988). "Freedom of the Will and the Concept of a Person." In *The Importance of What We Care About*, 11-25. New York: Cambridge University Press.
- . (1999). "On Caring." In *Necessity, Volition, and Love*, 155-180. Cambridge: Cambridge University Press.
- . (2006). *Taking Ourselves Seriously and Getting It Right*. Stanford, California: Stanford University Press.
- Fried, Charles. (1970). *An Anatomy of Values*. Cambridge, Massachusetts: Harvard University Press.
- Frost, Robert. (1913). "Reluctance." In *A Boy's Will*. New York: Henry Holt and Company.
- Gardner, John. (2015). "The Way Things Used to Be." Unpublished draft. (Cited with the author's permission, given on May 21, 2015.)
- Halbertal, Moshe. (2012). *On Sacrifice*. Princeton, New Jersey: Princeton University Press.
- Hampshire, Stuart. (1989). *Innocence and Experience*. Cambridge, Massachusetts: Harvard University Press.
- Harman, Elizabeth. (2009). "'I'll Be Glad I Did It' Reasoning and the Significance of Future Desires." *Philosophical Perspectives* 23, 177-199.
- Hieronymi, Pamela. (2005). "The Wrong Kind of Reason." *The Journal of Philosophy* 102, 437-457.
- . (2013). "The Use of Reasons in Thought (and the Use of Earmarks in Arguments)." *Ethics* 124, 114-127.
- Keller, Simon. (2013). *Partiality*. Princeton, New Jersey: Princeton University Press.

Khoury, Jack. (2012). "High Court Rejects Petition against Israel's Controversial 'Nakba Law'." *Haaretz*, January 5, 2012.

Kolodny, Niko. (2003). "Love as Valuing a Relationship." *Philosophical Review* 112, 135-189.

—. (2010). "Which Relationships Justify Partiality? The Case of Parents and Children." *Philosophy and Public Affairs* 38, 37-75.

—. (2011). "Aims as Reasons." In *Reasons and Recognition: Essays on the Philosophy of T. M. Scanlon*, ed. Samuel Freeman, Rahul Kumar, and R. Jay Wallace, 43-78. New York: Oxford University Press.

Korsgaard, Christine. M. (1983). "Two Distinctions in Goodness." *The Philosophical Review* 92, 169-195.

—. (1993). "The Reasons We Can Share." *Social Philosophy and Policy* 10, 24-51. Republished in Korsgaard, Christine. (1996). *Creating the Kingdom of Ends*. New York: Cambridge University Press. 275-310.

—. (1996). *The Sources of Normativity*. Cambridge: Cambridge University Press.

—. (1999). "The General Point of View: Love and Moral Approval in Hume's Ethics." *Hume Studies*, 3-42.

—. (2008). "Acting for a Reason." In *The Constitution of Agency*, 207-232. New York: Oxford University Press.

Lamb, Roger E. (1997). "Love and Rationality." In *Love Analyzed*, ed. Roger E. Lamb, 23-48. Boulder, Colorado: Westview Press.

Moran, Richard. (2012). "Kant, Proust, and the Appeal of Beauty." *Critical Inquiry* 38, 298-329.

Morgenbesser, Sidney; Ullmann-Margalit, Edna. (1977). "Picking and Choosing." *Social Research* 44, 757-785.

Morris, Benny. (2004). *The Birth of The Palestinian Refugee Problem Revisited*. New

- York: Cambridge University Press.
- Murphy, Liam. (2000). *Moral Demands in Nonideal Theory*. Oxford: Oxford University Press.
- Musil, Robert. (1986). "The Perfecting of a Love." In *Selected Writings: Robert Musil: Young Torless, Three Women, The Perfecting of a Love, and Other Writings*, ed. Burton Pike, 179-222. New York: The Continuum Publishing Company.
- Nagel, Thomas. (1970). *The Possibility of Altruism*. Princeton, New Jersey: Princeton University Press.
- . (1986). *The View from Nowhere*. New York: Oxford University Press.
- . (2014). "An Invitation to Hand-Wringing." *London Review of Books* 36, 25-26.
- Nietzsche, Friedrich. (1986). *Human, All Too Human*, tr. R. J. Hollingdale. New York: Cambridge University Press.
- Overvold, Mark. (1980). "Self-Interest and the Concept of Self-Sacrifice." *Canadian Journal of Philosophy* 10, 105-118.
- Parfit, Derek. (1984). *Reasons and Persons*. Oxford: Clarendon Press.
- . (2011). *On What Matters: Volume One*. Oxford: Oxford University Press.
- Proust, Marcel. (1930/1992). *In Search of Lost Time, Vol. II, Within a Budding Grove*, tr. C. K. Scott Moncrieff, Terence Kilmartin, and D. J. Enright, New York: Random House.
- . (1938/1999). *In Search of Lost Time, Vol. IV, Sodom and Gomorrah*, tr. C. K. Scott Moncrieff, Terence Kilmartin, and D. J. Enright, New York: Random House.
- Railton, Peter. (1984). "Alienation, Consequentialism, and the Demands of Morality." *Philosophy and Public Affairs* 13, 134-171.
- Rawls, John. (1971/1999). *A Theory of Justice*. Cambridge, Massachusetts: Harvard

- University Press.
- . (1993/1996). *Political Liberalism*. Cambridge, Massachusetts: Harvard University Press.
- Raz, Joseph. (1989). “Liberating Duties.” *Law and Philosophy* 8, 3-21.
- . (1999a). “Incommensurability and Agency.” In *Engaging Reason*, 46-66. Oxford: Oxford University Press.
- . (1999b). “Explaining Normativity: On Rationality and the Justification of Reason.” In *Engaging Reason*, 67-89. Oxford: Oxford University Press.
- . (1999c). “The Central Conflict: Morality and Self-Interest.” In *Engaging Reason*, 303-332. Oxford: Oxford University Press.
- . (2001). *Value, Respect, and Attachment*. Cambridge: Cambridge University Press.
- Regan, Donald. (1997). “Value, Comparability, and Choice.” In *Incommensurability, Incomparability, and Practical Reason*, ed. Ruth Chang, 129-150. Cambridge, Massachusetts: Harvard University Press.
- Scanlon, T. M. (1998). *What We Owe to Each Other*. Cambridge, Massachusetts: Harvard University Press.
- . (2004). “Reasons: A Puzzling Duality?” In *Reason and Value: Themes from the Moral Philosophy of Joseph Raz*, ed. R. Jay Wallace, Philip Pettit, Samuel Scheffler, and Michael Smith, 231-246. Oxford: Clarendon Press.
- . (2005). “When Does Equality Matter?” (unpublished).
- . (2014). *Being Realistic about Reasons*. New York: Oxford University Press.
- Scheffler, Samuel. (1994). *The Rejection of Consequentialism*. Oxford: Clarendon Press.
- . (2001). *Boundaries and Allegiances*. Oxford: Oxford University Press.

- . (2010a). “Valuing.” In *Equality and Tradition*, 15-40. New York: Oxford University Press.
- . (2010b). “Morality and Reasonable Partiality.” In *Equality and Tradition*, 41-75. New York: Oxford University Press.
- Schroeder, Mark. (2007). *Slaves of Passions*. New York: Oxford University Press.
- . (2012). “The Ubiquity of State-Given Reasons.” *Ethics* 122, 457-488.
- Setiya, Kieran. (2014). “Love and the Value of a Life.” *Philosophical Review* 123, 251-280.
- Stocker, Michael. (1976). “The Schizophrenia of Modern Ethical Theories.” *Journal of Philosophy* 73, 453-466.
- Strand, Mark. (2014). *Collected Poems*. New York: Random House, LLC.
- Stroud, Sarah. (2006). “Epistemic Partiality in Friendship.” *Ethics* 116, 498-524.
- Ullmann-Margalit, Edna; Morgenbesser, Sidney. (1977). “Picking and Choosing.” *Social Research* 44, 757-785.
- Velleman, J. David. (1999). “Love as a Moral Emotion.” *Ethics* 109, 338-374.
- Wallace, R. Jay. (2013). *The View from Here*. New York: Oxford University Press.
- Williams, Bernard. (1981a). “Persons, Character and Morality.” In *Moral Luck*, 1-19. Cambridge: Cambridge University Press.
- . (1981b). “Moral Luck.” In *Moral Luck*, 20-39. Cambridge: Cambridge University Press.
- Winterbotham, F. W. (1974). *The Ultra Secret*. London: Wiedenfeld and Nicolson.
- Wolf, Susan. (1992). “Morality and Partiality.” *Philosophical Perspectives* 6, 243-259.

—. (2010). *Meaning in Life and Why It Matters*. Princeton, New Jersey: Princeton University Press.