Beyond Reasons

Stephen Marrone

Charlottesville, VA

B.A. Philosophy, The University of Chicago, 2014

M.A. Philosophy, Georgia State University, 2018

A Dissertation Presented to the Graduate Faculty of the University of Virginia in Candidacy for the Degree of Doctor of Philosophy

Corcoran Department of Philosophy

University of Virginia

April, 2023

For my Mom, who would think it is neat that I am a philosopher

For my Dad, who still thinks that I am a psychologist

Acknowledgements

I am grateful to Talbot Brewer for his patient guidance over the last five years. Without his endless energy to read, revise, and meet with me again and again, this project would have never been formed. Finding the Aristotelian mean as an advisor is no small feat. But throughout my tenure in Charlottesville, Tal has offered guidance without dominance, encouragement without delusion and through his boundless energy for philosophical discussion, breathed life into my ailing confidence.

I am also grateful to Rebecca Stangl and Elizabeth Barnes for their thoughtful contributions to this project. As passionate teachers, careful readers and caring interlocutors, each have inspired me to be a more attentive teacher, scholar and person. I, and this project, have benefitted immensely from their help.

I am most grateful to my partner Emma. From offering friendship in some of the darkest moments of my life, to enduring years apart, to changing careers and finding new ways to make our life work, she has offered an unwavering support that could never be measured. Without her steady patience for dad jokes, passion for the work of making a better world and an absolutely unrelenting intolerance for bullshit, I would not believe that I have anything to say.

Finally, I place here some love for my dog Ginny. Learning to care for a being of another species is a unique education. While a series of world-altering catastrophes unfolded during the course of this project, Ginny offered steady structure to my days, instruction on proper squirrel stalking technique, and above all a world-class commitment to hydration, stretching and getting enough sleep. One could not ask for a better canine mentor.

Table of Contents

Introduction……..5

Chapter 1: Reasons and the Limits of Genericity……..18

Chapter 2: Reasons, Individuation and Practical Thinking in Dancy's Particularism……...58

Chapter 3: Rethinking Practical Thinking…….…….……...86

Chapter 4: Practical Thinking as Loving Attention...……..121

Conclusion……...159

References……167

Introduction

The sun was setting over the Shenandoah valley and I still had a *long* way to go. My knee was bruised, my lungs were burning and the cold of the last river crossing had turned my feet into stone. As I ran through the Blue Ridge Mountains, winding down through the leaf-covered forest trails, across valley streams and back up rocky ridges to the road, I had few articulable thoughts. When the hazy-blue daylight settled into dusk and I crested the final peak, I felt even worse. I was sore, wet, dehydrated, sunburned and ravenously hungry. And yet, as I limped down the last bit of trail to arrive right where I began, I felt grateful to have apprehended something valuable. When I stumbled home, stiff, cold and eating everything in sight, my wife asked "why do you do this to yourself?" I had no satisfying answer for her. I still don't.

What explains my eccentric choice of exercise? Philosophers typically attempt to answer this type of challenge by appealing to *reasons*. What explains my *action* (running) are considerations that count in favor. In this case, I might point to the stunning views, the company of friends, the benefits of exercise or even the peaceful calm of a mountain trail as considerations that count in favor of my endeavor. These are *generic reasons*: considerations whose generic description most any rational agent can be expected to find intelligible, regardless of whether they have undergone the particular experience themselves. Generic reasons are a common currency in philosophical models of practical thought. In fact, many philosophers take the articulation and comparison of generic reasons to be the central event in practical thinking in both individual deliberation and interpersonal justification.

However, attempting to capture the value of my activity this way raises more questions than it answers. It is surely right that my decision to run (and keep running) was an intentional action—nobody was forcing me. I decided (repeatedly, and sometimes with great reluctance) to

put one foot in front of the other. It is equally true that, in retrospect, I affirm those generic reasons listed above: it *is* beautiful in the mountains; the views *are* stunning; exercise *does* feel good. And yet, none of those generic reasons for running capture the value that I was responding to in the activity, the good I found from within the unfolding work of moving through the mountains. To run for hours, at great cost for little or no material gain, only to end up right back where I began, might seem positively irrational.[1] At the very least, it cries out for explanation.

But supplying a reasons-based explanation is tricky; on the one hand, it seems wrong to suggest that I acted on the strength of generic reasons. Though I appreciate the mountain views, the exercise and the quiet, none of these were present to mind during the activity. On the other hand, it seems equally strange to suggest that I ran all that way with *no intelligible motivation at all*. I want to suggest that this puzzle is the result of an underexplored oversight in philosophical accounts of practical thinking. Where most prominent philosophers take reasons for action to be central to practical thought, the way we apprehend the value of activities like long distance running resists formulation in those terms. This, I will argue, is because the value of an activity like long-distance running is only fully appreciable from within a prolonged attempt to do it. Absent any relevant practical experience, it is difficult to get in view what is worth going for in such a painful exercise in terms that any rational agent might be expected to find intelligible. Since the standard philosophical methodology for modeling how thought moves into action takes the concept of *reasons* to cover the extent of the possibilities for practical thinking, philosophy is ill-equipped to explain the good of embarking on this kind of activity. It seems that we have run

[1] Thomas Nagel (1986) infamously said that he finds marathon running completely unintelligible. He was joking. Nonetheless, his quip lights up an important point here: many people could affirm the truth of the propositions that might count in favor of such an activity (e.g., scenic views, fitness goals, etc.) without having any understanding of or desire for the good of the activity.

into an instance of genuine practical thinking where the value of its object *resists* formulation in terms of reasons for action. When prompted to explain what I saw worth going for in my extended run, *talk of reasons runs out*.

What accounts for this discrepancy? Cora Diamond provides a helpful starting point. She calls philosophy's characteristic tendency to misrepresent reality *deflection*. Deflection is "what happens when we are moved from the appreciation, or attempt at appreciation of a difficulty of reality to a philosophical or moral problem apparently in the vicinity."[2] In the case above, the temptation is to mistake the difficulty of apprehending the value of running as a difficulty about articulating the right reasons for action. These difficulties are distinct.

> The hardness there, in philosophical argumentation, is not the hardness of appreciating or trying to appreciate a difficulty of reality. In the latter case, the difficulty lies in the apparent resistance by reality to one's ordinary mode of life, including one's ordinary modes of thinking: to appreciate the difficulty is to feel oneself being shouldered out of how one thinks, how one is apparently supposed to think, or to have a sense of the inability of thought to encompass what it is attempting to reach. (Diamond 2003, 12)

A *difficulty of reality* marks experiences where we encounter something in reality which resists our way of thinking about it. The *difficulty of philosophy* is the difficulty of bringing the difficulty of reality into philosophical focus. Deflection is what happens when philosophers imperceptibly mistake the difficulty of appreciating some part of reality for a problem peculiar to philosophical argumentation. The philosopher *rethinks* the problem in the language of philosophical discourse.

[2] Diamond (2003), 12.

As a result, an understanding of that reality, perhaps a unique window onto human life and struggle, is diminished or distorted when reimagined in philosophical argument.[3]

I think that philosophy's reliance on talk of *reasons* can result in deflection. In our case, deflection away from a difficulty of reality (appreciating the significance of running) and toward a philosophical problem (articulating reasons for acting) is rooted in a common conception of the discipline itself. The way contemporary philosophers in the analytic tradition typically understand the life of the mind essentially involves providing reasons in argument. As philosophers, we are always tempted to deflect toward reasons, in part, because we take philosophical work to be centrally concerned with reasoning toward action or beliefs. Because reasons are so clearly connected with the self-image of the philosopher, it is very difficult *not* to deflect in this way.

So, my initial claim is that reasons can, in some cases, deflect us away from the phenomenon of individual valuing.[4] This is a likely spot for philosophical deflection because, once again, reasons are central to what we do: in the most general sense, a philosopher's job is to articulate a reasoned basis for others to join in our argumentative conclusions. Seen in this way, philosophy is finely tuned for tackling certain problems. Through its study we can learn how to discuss hard issues, we can learn what constitutes a good argument, various ways they can be distorted and so on. This can be quite helpful. But such a fine-tuned conception of the field can

[3] *Ibid*., 11-12. Sometimes, the fact that the original problem has been recast in terms exclusive to philosophy can itself be evidence of distortion. Diamond's example is Peter Singer's response to the novelist J.M. Coetzee's Tanner lectures (later published as *The Lives of Animals*). Where Singer takes Coetzee's lectures to be ethical arguments about how we ought to treat animals in general, Diamond rightly points out that this is to mistake the difficulty presented. Insofar as Singer (and others) are narrowly focused on Coetzee's words as reasons in philosophical argument, casting the examples as general claims rather than poetry, Singer misses the real difficulty at hand in Coetzee's lectures: understanding the kind of animals we are.

[4] It is interesting to note here Diamond's note on the sometimes underestimated differences between individuals: "the things we take [to be difficult to wrap our minds around] may simply not, to others, present any difficulty—of being hard or impossible or agonizing to get one's mind round." (2003, 3) As we move through the following chapters, the difference between the conceptual possibilities across individuals will become increasingly significant.

also pose a challenge. Diamond's insight is that the tools of philosophy and the ways of thinking that such tools cultivate (e.g., logic, argument, thought experiments, a desire for clarity, etc.) can deflect us from the nature of the very difficulties we philosophers are attempting to bring into view. In our starting case, the philosophical difficulty emerges as a search for a suitable description of which *reasons* long distance runners respond to. However, the difficulty of understanding the value of this activity looks to be another thing entirely.

But before condemning reasons to the flames, let us lay some groundwork. What *are* reasons? Where do they come from and why do philosophers take them to be the center and relevant extent of practical thought? Even a cursory glance at the literature is likely to yield a headache. Ask *whence* reasons arise and you find a veritable river of writing.[5] Ask what they *comprise* and you find an ocean.[6] Nevertheless, in order to begin assessing what work reasons can and cannot do, we must first have a general gloss.

In philosophy, reasons are typically divided into two kinds: normative and motivating.[7] A *normative reason* is a consideration that counts in favor of an action or conclusion. One way of understanding this claim is in terms of justification: a normative reason makes it right for someone to act in a certain way. For example, the fact that my dog has ingested a potentially lethal amount of chocolate makes it right for me to get her stomach pumped. While this may be

[5] Accounts of the origin of normative reasons often focus on the *relation* between motivating and normative reasons. Some, for example, argue for a *causal* connection. See e.g., Wedgwood (2006) and Turri (2011). Others argue for a *competence* view, suggesting that normative reasons are evidence of a certain competence of the agent. This view has its origin in Aristotle (NE 1105a–1105b). In contrast, Susanne Mantel (2018) provides an excellent overview of the field while advancing a *correspondence* view of the relation between motivating and normative reasons.

[6] See, e.g., Smith (1994), Scanlon (1998; 2014), Dancy (2000; 2004), Stoutland (2001), and Alvarez (2010). In a recent book, Charles Larmore is willing to go one step further, defending a form of "Platonism about reasons" according to which reasons constitute "a third dimension of reality, neither physical nor psychological" (Larmore 2021, 7–8). I will discuss Larmore's view directly in chapter three.

[7] This is a modern distinction, though it can be traced at least as far back as Francis Hutchinson (1971). While it is possible to question its legitimacy, for our purposes, we need only note that many contemporary philosophers accept and defend this way of dividing up the landscape of reasons. See, for instance, Raz (1975); Smith (1994); Parfit (1997); and Dancy (2000).

painful for the dog, the fact that her life may be endangered counts (decisively) in favor of taking her to the vet. Though I can fail to be motivated by such concerns (e.g., if I am cruel, ignorant or otherwise mistaken) normative reasons provide the justificatory grounds for acting as I did.

A *motivating reason* is a consideration that, *in the eyes of the agent*, counts in favor of doing some action. When I am motivated by a reason, I act in its light such that the consideration is a premise in my practical reasoning, whose result is (or can be) an action. For example, when my dog whines, I take the behavior to count in favor of giving her a treat. Whether or not I am mistaken that her behavior really does count in favor of providing a reward (e.g., she may have trained *me* to regularly provide treats when she does not really need them) whining is a motivating reason for me. Where a normative reason *in fact* counts in favor of some action, a motivating reason is a consideration that the agent *takes* to be normative.

Motivating reasons can also be used to *explain* what one does. As I type these words, I act in light of the consideration that moving my hands over certain keys in a certain order will render my thoughts in writing, hopefully with clarity and precision. As such, a motivating reason can explain my movements by placing these acts under a description that specifies what I took to be worth going for. The function of this type of reason is to explain what I thought I was doing.[8]

So, reasons can be called to justify, motivate or explain. What observations can we draw from such a basic assessment? First, divergent views about reasons seem to share a common referential structure.[9] Philosophers of many stripes assume a relatively straightforward account of

[8] Dividing the functions of practical reasons (reasons concerning what we ought to do) this way does not necessarily entail that motivating and normative reasons are different in kind. Instead, some have suggested that what marks this distinction is simply what kinds of questions we can rightly ask of each. Jonathan Dancy makes this point clearly. On his view, to call a reason “motivating” is merely to "issue a reminder that the focus of our attention is on matters of motivation, for the moment." Similarly, he claims, calling a reason “normative” is "stressing that we are currently thinking about whether it is a good reason, one that favours acting in the way proposed." (Dancy 2000, 2–3).
[9] This structure is referred to "in virtue of" concept ordering, and informs the central argument of chapter four, section two.

the relation between reasons and some neutrally describable set of facts. Some hold that what one has reason to do depends upon which propositions are true about the world.[10] Others hold that one's reasons depend upon facts about the agent's desires.[11] While there is much disagreement about the origin and nature of reasons, a common referential structure can be seen. Some features (of agent themselves or the world) are taken to exist beneath and ground which reasons we have. These features fix when the term *reason* properly applies.

Another observation is that normative reasons are used as an interpersonal currency. We ask for and formulate reasons to explain and justify ourselves to each other. This means that we are in need of standards for what counts as a reason. Different people who share different values, points of view and epistemic limitations will take different positions on which considerations do and do not count as reasons to do something. With a desire to rescue reasons from the threat of cultural relativity, philosophers often point to a conception of what we are aiming to do when formulating reasons—a genuine reason must be a consideration that anyone sufficiently situated can access, otherwise it could not be offered to anyone, anytime in a universal marketplace of moral exchange.[12]

Such an approach lends itself to the metaphor of *weighing*.[13] Philosophers like to speak of weighing reasons, comparing their relative *strength* or *grip* on us; so much so that some have made it the basis of their account of practical reasoning.[14] Weight is a helpful concept here because it can be abstracted from objects without knowing those objects; we can know the

---

[10] See Raz (1975) and Scanlon (1998), though they have different accounts of the nature of these facts.
[11]This family of views also contain quite a variety: compare, for instance, Williams (1981) with Schroeder (2022).
[12] This conception of reasons is discussed in chapter one, section two and is referred to as the *formal* approach.
[13] The metaphor of weighing will be revisited in chapter one, section one as well as the end of chapter four.
[14] In chapter one, section one I refer to this model of practical thought as the *naive view* because, as we will see, there are cases where it seems that this model fails to track the phenomenology of practical thought: we often compare competing claims on us without anything that looks like weighing.

weight of a gemstone or distant planet without ever encountering them. As a tool of comparison, weight provides a common scale and commensurate unit of measurement across which we can assess apples and oranges.

Digging further, if reasons are to serve either function (explanation or justification) reasons must also be articulable in words.[15] If I am to explain or justify my action, I need to do so in words that others can be reasonably expected to understand in their full force and meaning. If I fail to say anything intelligible, then it is trivially true that I have not explained or justified anything. Reasons are thus dependent on a minimal standard of articulacy in order to serve either of their central purposes. Once again, the generality of this assumption makes it difficult to see: how else are we to explain or justify ourselves if not in words? But as we will see in chapter one, this standard takes for granted an articulacy that is not always possible.

Finally, if reasons are to serve their dual function in ethics, they must not only be articulated in words, but also rely on a broadly shared understanding.[16] If you ask why my fingers are flitting over these plastic keys, and I offer the explanation that I am writing a philosophy paper, whether that counts as a sufficient explanation relies on our shared understanding of what a keyboard is, what a philosophy paper is, and how those two things are connected through the act of typing. While someone unfamiliar with computing technology may understand the words I offer, without some basic cultural, historical or social touchstones, such as a familiarity with electronic communication and its uses, they will fail to understand my movements as intentional action.[17]

---

[15] I take up this topic in chapter one, sections three and four.

[16] I discuss the role of cultural and historical circumstances and their influence on the imagination most directly in chapter three.

[17] Anscombe (1957) playfully makes this point by describing a confused neighbor watching a man carefully arranging saucers of mud.

With normative reasons, the relevant shared understanding involves what ends are, in general, worth going for. Consider an analogy. Just as the value of currency is dependent upon a tacit, shared understanding of the equivalent good(s) it can be exchanged for, whether a consideration counts as a reason (and what importance it is assigned) depends upon a shared understanding of its normative significance. Where there is disagreement about the good or the right, there will be disagreement about whether and to what extent a consideration counts as a normative reason. At the crossroads of such disagreement is where our investigation begins.

This dissertation will argue that the dominant conception of reasons is not adequate for its central task. Specifically, it argues that while reasons are taken to be the locus of practical thought, there exists a familiar range of activities in everyday life whose full appreciation escapes meaningful formulation in terms of reasons. As a result, the significance of these activities is systematically discounted or in some cases ignored by the dominant philosophical assumptions about what counts as a reason.

Chapter one begins by examining two dominant approaches to understanding generic reasons in moral philosophy. It argues that because these approaches depend on the sufficiency of generic reasons for modeling practical thought, they end up excluding an important class of human experiences, activities and relationships from fair consideration. This class is broad, ranging from studying philosophy to appreciating art to being in love. Following Aristotle, I call members of this class *energeia*. While obviously central to moral deliberation, the value of an energeia essentially resists being heard in the language of public exchange. What is worth going for in our most deeply held pursuits cannot always be made accessible to the wider public, but must rather be unfolded oneself. Since both approaches rely on generic description to capture the nature and importance of candidate reasons, both systematically misrepresent energeia, the full

appreciation of which depends upon direct experience. The difficulty of representing this way of valuing in terms of generic reasons, and indeed getting this difficulty in philosophical view at all, marks it as an instance of deflection.

This presents serious problems for philosophical ethics. If there are important human pursuits whose value cannot, in principle, be fairly represented by the way theorists think about reasons, then not only will the value of such endeavors be improperly weighted in moral deliberation, they might be mistakenly excluded entirely from the domain of morally relevant reasons, or perhaps reasons more generally. More worrying still, their exclusion could not possibly be justified. On the one hand, if such considerations cannot meet some formal condition for what counts as a reason because they cannot even be articulated in that form, then the basis for their exclusion cannot be given in that language either. What cannot be recognized cannot be dismissed. On the other hand, if such considerations are recognized as such, but deemed morally irrelevant, then a central swath of human life is left out of theoretical accounts of moral deliberation. That this problem of exclusion is common to both dominant ways of representing reasons suggests that the heart of the issue is not the content of these approaches, but rather the philosophical methodology that relies heavily on reasons to model practical thinking. Put another way, chapter one reveals a methodological limitation on what reasons can do.

Chapter two considers a third alternative beyond the formal and substantive approaches to delineating reasons. Specifically, it outlines why particularism about reasons could be a natural solution to our worries from chapter one. Focusing on the work of Jonathan Dancy, it then argues that despite this appeal, particularism's most prominent proponent fails to harness that potential because his view shares some foundational assumptions about the function of reasons in practical thinking with the generalist views we have already discussed.

What is revealed is that the debate between Dancy's particularism and generalism only begins after another, more fundamental set of methodological assumptions about reasons in practical thinking are already in place. Signing on for these assumptions has the added consequence of missing the phenomena that form the basis of a strong, yet overlooked objection to generalism. While Dancy's view offers a key revision (or rejection) to the methodology of mainstream moral theorizing by embracing the uniqueness of individual experience, being sensitive to contextual details and rejecting the universalizing impulse of most analytic moral theorizing, I conclude that his view does not go far enough.

Chapter three draws on the work of Iris Murdoch to reconstruct an alternative way of understanding the phenomenon of valuing energeia by rebuilding the concepts of attention, imagination, love and practical thought. In particular, it explores the early philosophical work of Iris Murdoch, asking what resources her account of practical thinking and moral deliberation might offer. The chapter then interprets Murdoch's work on love, attention, imagination and the phenomenal experience of concept acquisition to illustrate how philosophical inquiry might better capture the modes of valuing we discussed in chapter one than the approaches that are dominant today. By considering and then sharpening these terms, this chapter lays the groundwork for my own positive account of practical thought in the final chapter.

Chapter four caps the project with a positive account of practical thinking called *loving attention*. On its own, loving attention attempts to depict the appreciation of value in the central phenomena that began our investigation. As an amendment to a particularist view of moral deliberation, it kicks in where reasons talk runs out, contributing a richer understanding of the elements of practical thought in some familiar cases of valuing energeia, including concept acquisition, imagination, and attention. I show that doing so requires reconceptualizing the scope

and point of practical thought with an eye toward accommodating a wider range of activities and experiences in human life. After answering objections, I conclude by illustrating two important benefits of this view. First, that it better accounts for the unity and depth of our way of valuing energeia than the dominant reasons-focused philosophical methodology. Second, that it represents an underexplored blurring of the boundaries between theory and practice that raises doubts about the clarity and distinctness of the category of specifically *practical* thinking itself.

Overall, this project will deliver three contributions to the field. First, it provides a clear statement of the under-explored tension between particular valuing and philosophy's commitment to genericity. To extend an economic metaphor, if reasons are the currency and values the goods, then there are goods in the marketplace of life that cannot, in principle, be assigned an accurate price in that currency. And if reasons are taken to be the entirety of what makes an action right or wrong, yet they do not capture the whole of what it is we value in some important and central cases in life, then there is a serious and unacknowledged limitation on the normative force of the substantive conclusions of philosophical theories in ethics. Second, in uncovering this shortcoming of philosophy's commitment to generic representation of particular values—and the subsequent limitations embedded in its background assumptions about practical thought—this discussion offers a potentially powerful argument against the use of general principles. Finally, it will outline a new account of practical thinking that makes better sense of the phenomenon of individual valuing in cases where its representation in terms of reasons falls short.

One last preliminary note. Though I will show that reasons can deflect in the ways outlined above, the structure of my argument will make it difficult to say just what reasons deflect us *from*.[18] Since I will contend that the language of reasons is insufficient to fully appreciate the

[18] I take up this problem again in chapter one, section five.

value of some important activities in human life, a perfect specification of what is missed in those same terms (using generic reasons, examples and scenarios) will not be possible; indeed, success in doing so would undermine my overall argument. What I offer instead can be seen as an extended invitation for the reader to meditate upon the value of such activities in your own life. To feel the force of my overall argument, I will ask you to think alongside me.

Chapter 1: Reasons and the Limits of Genericity

I

Late one April evening, I sat puzzling over my *adjusted gross income* with sweat on my brow. As I feverishly worked to the wire to get my taxes done on time, my wife suddenly burst through the door. With a gasp of exasperation, she threw her bag on the ground and collapsed on the sofa. After a moment of pause, I dropped my W2, then hurried over to make a cup of tea and help her unwind from a long day at work—taxes be damned.

Suppose I acted just as common-sense morality prescribes: when a loved one needs support, my taxes can wait. What explains why this is so? Mainstream philosophy offers a confident—and I think ultimately unsatisfying—range of explanations. On one view, the rightness of throwing my taxes aside is explained by reference to the overall balance of my practical reasons. Practical reasons are the considerations that speak for and against alternative courses of action. Deliberation about what to do begins with an attempt to gather up and weigh one's practical reasons. Call this the *naive view* of practical moral reasoning.

On the one hand sits my unfinished obligation to the IRS, whose weight is increased in inverse proportion to the time left before I am assessed a late fee. To this, one might add the need for continuity and focus in order to accomplish the task at hand. The fact that I have already been at this work for hours, taken together with my need to remember where I am in the process, provides further reason, one might think, to ignore my wife's sigh and tidy up my work. These, then, are the considerations that count in favor of finishing my taxes.

On the other side of the scale sits the obvious fatigue of someone I love. Given our relationship and my standing commitment to support it, the fact of her arrival, together with her

state of mind, presents a strong consideration that counts in favor of putting down my forms. On this view of practical reasoning, what happened in this case was that I simply assessed all these considerations, applied what I took to be their proper weight, and once I determined which side tipped the scale, I made a choice.

Such a model works well in some cases. But weighing is a poor analogy for what has happened here. Have I truly *weighed* the urgency of my taxes against the needs of someone I love? After all, my affection for my wife bears little resemblance to the anxious nagging that flows from my taxes. Where the former is one of the deepest, most important goods in my life, the other is a surface level annoyance at most.

The defender of the naive view might reply that surely these are very different goods, but that is why the considerations that flow from my wife and my taxes carry differing *weights*. On this view, we account for these obvious differences by assigning varying weights to their respective resulting considerations. What explains my behavior, then, is simply that the needs of my wife outweigh the need to file my taxes. Job done.[19]

Many philosophers reject the naive view. Some argue that, in cases of crisis or urgent conflict, a tendency to pause and weigh one's options *itself* reveals a deficiency in character.[20] Others have argued that proper recognition of some reasons *silence (*rather than outweigh) conflicting reasons.[21] Still others contend that weighing is sometimes unnecessary because

[19] Following this thread suggests that the metaphor of weight, ubiquitous as any in ethics, is misleading. To weigh is to abstract away from the properties of a particular object; its work is to facilitate the comparison of unlike things. Importantly, one can know how much something weighs without knowing anything else about it. The question I want to pose here is whether this is an apt or informative metaphor to serve as the center philosophical models of practical thought.

[20] Williams (1981) characterizes this as "one thought too many." Wolf (2012) offers a helpful reinterpretation of this oft-repeated phrase. But it is worth pointing out that most commentators (including Wolf) seem to miss that Williams is not merely concerned with *when* deliberation takes place (be it delaying rescuing a drowning spouse in the moment or wrestling to justify one's action after the fact) but whether practical deliberation in the weighing of reasons is an appropriate response at all.

[21] This is McDowell's (1979) term.

occupying a certain role (e.g., being a good partner) brings with it a demand to bracket or reorder the reason-giving force of some of one's personal concerns (e.g., to finish my taxes) that may conflict with the duties of the role.[22] We ought, then, to reject the naive view: in its reliance on the weighing analogy, the naive view fails to track the experience of practical reasoning. Perhaps my reasons to finish my taxes were not outweighed, but rather (rightly) *ignored*, *silenced* or simply made *irrelevant*.

However, supplanting the naive view is harder than it may appear. Mainstream moral theories rely heavily on the articulation of generic reasons for action to model what is worth going for in a given situation. A *generic reason* is a consideration that is fit for understanding by any rational agent regardless of their practical experience. For ease of discussion, we can say that this reliance on the formulation of generic reasons constitutes a *standard view* of practical moral reasoning. Moral theories that fall under the standard view are marked by a common methodological strategy for modeling human valuing: they take the language of generic reasons as perfectly adequate for representing the value of anything worth going for, or if not, holding that those goods whose value resists formulation in those terms can be rightly left outside the domain of morality. The reliance of the standard view on generic reasons for action marks a tendency of moral theorists to collapse the value of radically different human goods (and the variety of ways we come to value those goods) into a single form.

In this chapter, my first claim will be that we should amend the standard view because its central nexus, the language of generic reasons for action, fails to fairly represent certain ways of valuing that can be crucial to making decisions. To make this case, I will detail examples of

[22] Scanlon (1998) uses the example of croquet to highlight this type of situation, but the point could be extended beyond sports to other formal or informal roles that demand the reordering of the reason-giving force of personal considerations.

activities, experiences and relationships whose value can only be appreciated through the first-person attempt to bring their worth into view. Since such values are only unveiled in and through practical experience, what is worth going for in these kinds of activities is not available for general representation in terms of generic reasons. This opens a familiar schism between the *representation* of values in moral theories on the one hand and the *experience* of personal valuing on the other.

Some proponents of the standard view might concede this point, agreeing that this way of valuing is difficult to represent as reasons fit for general evaluation, yet deny that this ought to worry us.[23] The central work of a moral theory, they might argue, is to set principled limits on what we can permissibly do, which is a decidedly *practical* aim. As such, it is not necessary to bring forth the specific details of the complexity of individual human valuing in order to get on with the business of normative ethics. We simply don't need to know the details of *how* individuals value their personal commitments in order to recognize and respect *that* they do.

Moreover, defenders of the standard view in ethics might say that theorists have solid, pragmatic reasons to avoid getting bogged down in the details of the phenomenology of individual valuing—in the end, impartial moral principles must take precedence. To avoid violence, theft or general political chaos we ought to, on the whole, avoid giving priority to the pursuit of personal projects when they conflict with our moral obligations. If so, then whatever differences exist between the way that individuals actually experience different values and how those values are represented in theory, it makes no practical difference—when the considerations that flow from personal valuing run up against our moral obligations, the path is clear.

[23] I will take up such views in detail in section five of this chapter.

But whether that is true depends upon having an adequate cognitive grasp on the categories of value that could potentially conflict with our moral obligations. My work is to make visible what is veiled by the standard view, whose conceptual boundaries are set by only what can be said and understood in terms of reasons in general. This means highlighting the inadequacy of generic reasons for capturing the nature of individual valuing. If we are to make philosophical sense of the swath of human life that I highlight below, we need a view of practical thought whose essential concepts move beyond generic reasons (and as I will argue in chapter two, beyond reasons, period).

My second claim in this chapter will be that this work is worth doing. Theories of normative ethics purport to offer us, among other things, a picture of how we ought to reason, how we ought to settle real conflicts in our lives. The terms philosophers offer for adjudication between practical alternatives are typically offered in the spirit of recommendation; normative moral theories are, at least in part, guides for how our own thinking should go—and that thinking is entirely in terms of reasons for action. To the extent that ethics does offer prescriptive ideals for conflict resolution, the restriction of its tools for such adjudication to only reasons, taken to be generic considerations fit for public appreciation, raises at least three concerns.

First, if all that is visible in a moral theory, indeed all that could *ever* be visible, about individual human beings is their capacity as agents and repositories of general or nonspecific value, this would be dehumanizing.[24] The cases I take up illustrate that the language of generic

[24] Michael Stocker (1976, 458-459) makes this point. He argues that moral theories fail by making it impossible for a person to "achieve the good in an integrated way." On his view, they do so by offering justifications for why we ought to uphold certain principles of right that we could not psychologically embody. What I have termed a schism he calls a *schizophrenia* between the reasons we have for action in theory and the motives we actually live by. Following a similar path, Bernard Williams (1973; 1981) calls the further imposition of the authority of the impartial moral perspective on the psyche of individuals an attack on one's *integrity*. I discuss these views in more detail in section four.

reasons for action contains only the conceptual tools to *name* some of the values that individuals pursue, not bring them into view—after all, we do not experience all value in the world entirely in terms of generic reasons.

If this is right, it raises a further worry: absent a method for fairly representing the wider variety of ways of valuing, moral theories begin to look suspiciously *incomplete*; some of what actually moves human beings is either made inarticulable or deemed morally irrelevant in theoretical models of that very thing.

Some might still shrug at these concerns, insisting once more that moral theories make little claim to represent the *whole* of life in all its complexity. Be that as it may, to admit this much raises a third problem: on what grounds can morality claim normative dominion over actions in general, which concern values it cannot bring into focus? If one functional task of normative ethics involves setting limits around what we can and cannot do, the legitimacy of *that* work rests on the assumption that such limit placement is informed by a fair representation of what considerations exist on either side. But if one already agrees that the philosophical tools on offer are not capable of fairly representing the whole field of relevant practical concerns, then the ability of such theories to settle disputes is thrown into question.[25]

The following section begins by outlining two dominant rationales for what I have called the standard view, whose assumptions about reasons work to exclude or minimize certain ways of valuing that are centrally important in our everyday lives. I suggest that what philosophers take to be the central locus of practical thought—what is fit for serving the role of reasons—*itself* embodies a commitment to some unacknowledged substantive claims: in order to serve as a *real*

[25] It leaves us, as Gaita (2004, 23) writes, with "little reason to believe that the academic practice of moral philosophy has the authority to determine the best style and method of thinking on moral matters, or even what the most serious problems are and how they should be characterised."

reason, a consideration must be like an ordinary object: in some sense *there* (be it given or constructed) for all competent, rational users of descriptive language to see. Progressively unfolding activities such as philosophical study, the appreciation of art or individual human beings, do not fit this mold. The point of such activities are not always articulable, not simply given by the world, but is rather created, in part, by the first-personal attempt to engage with them. If this inarticulacy is due, in part, to the limiting assumptions about practical reasoning that ground dominant approaches to moral theorizing, it will provide motivation for the account I build in chapters three and four—a better way of representing why we act in light of activities we can neither fully appreciate or explain through generic reasons.

II

There are two popular routes to the standard view. The first is to place substantive restrictions on which considerations are morally relevant. Call this the *substantive approach*. The dominance of this thinking stems, in part, from a widespread aim of moral theorists to affirm only those moral claims that can be justified in a public deliberative exchange. On the substantive approach, morally relevant reasons are delineated by asserting substantive claims about the scope and point of morality. David Gauthier's approach to reasons in moral justification has this structure. On his view, morality represents a rational constraint on what is in our interest to do. How this rational constraint is defined for each of us turns on the stipulation of a set of substantive claims about which considerations do, in fact, lead us to our own best interest. While there are plenty of other considerations that may be considered viable reasons, which can work outside and even against the constraints of morality, normative reasons fit for moral justification

are only those that fit within the structures of interaction that comprise moral life: maximizing rationality within agreed mutual constraint.[26]

For Gauthier, moral reasons must be impartial because that is what is required by the structure of our interactions with others. If moral reasons must be universal, it is because that must be the scope of practical moral demands. And if moral reasons must be generic, this is because moral principles are voluntary agreements made between determinate individuals in pursuit of their individual self-interests; they must be shared in order to exist, and they make claims upon us that we have grounds not to ignore in favor of our personal inclinations.[27] So conceived, this approach suggests that moral justification involves only considerations that are publicly intelligible, generically rendered in democratic deliberative exchange and seen from an impartial point of view.

The second approach places a formal condition on reasons in general: perhaps *all* reasons are morally relevant, but some putative reasons are merely apparent. If I have a genuine reason to do something, then anyone who finds themselves in a sufficiently similar situation also has this reason.[28] Call this the *formal approach.* Charles Larmore puts it this way; “the reasons we seek to discover are by definition universal: if they are valid for us, then they are valid for all under similar circumstances.”[29] This closely echoes T.M. Scanlon’s formulation: "whenever we make judgments about our own reasons, we are committed to claims about the reasons that other people have, or would have under certain circumstances."[30] Any consideration that does not fit

[26] Gauthier (2022), 3.
[27] Gauthier (1986), 9.
[28] While it may have its origin in Kant, this view has been assimilated by theorists of many stripes. Scanlon (1998; 2014) refers to it as the *universality of reason judgments*. Korsgaard (1996; 2009) calls it a *publicity requirement*. Hare (1963) refers to *universal reasons* and Sidgwick (1981) calls it the *principle of equity*.
[29] Larmore (2021), 30.
[30] Scanlon (1998), 74.

this form is not a genuine reason. Scanlon calls this claim the universality of reason judgments. Christine Korsgaard takes this approach as well. For her, following Kant, "[u]niversalizability is a condition on the form of a reason, and if a consideration doesn't meet this condition, then it is not merely outweighed—rather, it is not a reason at all."[31] On this view, taking something to be a reason for action just is to see it as sufficient grounds for its conclusion for anyone similarly situated. It is precisely this publicity of normative force that moves reasons into the territory of morality. More than a substantive moral demand, publicity in this sense is a rational constraint on what a reason must be. After all, if a reason only applies to me, it is hard to see how I could take it seriously as a reason for me at all. In fact, some have even thought this so obviously correct that it does not require sustained argument.[32]

I think that both approaches contain a serious oversight. Specifically, there exists a class of ethically relevant human endeavors that are essentially resistant to being formulated in the language of generic description that is required by both approaches. Grasping the value of these pursuits requires a sustained, first-personal engagement with the experience, activity or relationship that is, in principle, not available for understanding to all. Without this unique experience, appreciation of their distinctive value will be distorted or passed over entirely. The machinery put forth (above) to evaluate candidate reasons represent two different ways of excluding the good we are aiming at in certain endeavors. This is because both approaches rest

---

[31] Korsgaard (2009), 51.

[32] For instance, R.M. Hare (1963, 89-90) tells us that, "[w]hen we are trying, in a concrete case, to decide what we ought to do, what we are looking for [...] is an action to which we can commit ourselves [...] but which we are at the same time prepared to accept as exemplifying a principle of action to be prescribed for others in like circumstances." Scanlon (1998, 74) echoes this thought: "whenever we make judgments about our own reasons, we are committed to claims about the reasons other people have." He concludes that this "is not something that should be a matter of controversy [because it is] a formal consequence of the fact that taking something to be a reason for acting is [...] a judgment that takes certain considerations to be sufficient grounds for its conclusion." Peter Singer (2011) even goes so far as to claim that something like this view was held by everyone from Moses, Jesus, the Stoics, Hume, Kant, Bentham, Rawls, Habermas, and Jean-Paul Sartre.

on the assumption that it is always, in principle, possible to capture the value of all relevant human goods in the language of generically described reasons. The failure of that assumption, as I will demonstrate below, suggests that one or the other dominant approaches I have sketched above, is incomplete. As theorists, if our models of moral reasoning fail to bring these important goods onto the playing field of discussion in their full force and form, our conclusions will be correspondingly limited.

Let us examine this excluded class. When I took my first philosophy class, what good there is in being a philosopher was not (and could not have been) fully available to me. I was interested, but still lacked the essential experience to see through to the depth of what philosophical study could be. I could not have articulated those goods to others or even myself in terms of reasons because I did not yet have the requisite experience to *see* them. In order to fully grasp why anyone would spend their days plumbing the depths of dusty old philosophy books or sit through hours of technical seminars, one needs to first do some (or often quite a lot) of philosophy.[33]

One might reply that, surely, *some* appreciation of philosophy is not conditional on attempting to study it for oneself. For instance, to claim that philosophy is valuable is to endorse a whole range of reasons—each from differing perspectives—as legitimate.[34] One is that some people who have the relevant ability and opportunity have reason to take up philosophical study as a career and to devote their lives to it. Another might be that those who take up philosophy as

[33] MacIntyre (2018, 145) describes this process for a novice of classical music. "As she moves forward, she becomes able to ask questions that at first, she would not have understood, but only because she has become able to appreciate works that at first, she would not have been able to hear, let alone appreciate. She has learned to perceive what without this initiation she would never have perceived." A similar initiation process is required before some particular values can be appreciated in their full depth.

[34] Here I am following T.M. Scanlon's (1998, 90-91) discussion of the relation between values and reasons. His example is of scientific inquiry, but he is quite explicit that he thinks this structure can be applied to many values.

a career have reason to try to be *good* philosophers: to work hard, to choose lines of inquiry that are significant rather than those that are merely popular. A third might be that if philosophy is, in fact, valuable, we have reason to try and understand it and support those who engage in this endeavor with us. In short, if philosophy is valuable, there are a range of potential reasons for supporting it that need not depend upon first undergoing its study for oneself.

It is possible to have *some* understanding of the category of value that philosophy embodies—but is it possible to fully appreciate such value *for oneself* without first encountering instances of its kind? While we may have generic reasons to study or support philosophical endeavor, the good inherent in philosophical study, that which pulls us into its grip when we attempt to understand a difficult puzzle, cannot be fully understood absent some practical experience. This is Aristotle's point about his audience at the outset of the *Nicomachean Ethics*. In his view, the young are not an appropriate audience for political or ethical expertise because they lack precisely the practical experience of the actions that *constitute* political life.[35] Aristotle tells us that the appreciation of theoretical study of politics, art, or ethics is only really useful when one already has the raw materials of practical experience to draw upon and mold into form. In ethics, he thinks that in a given case the ultimate determination of right and wrong, all things considered, (*prohairesis*) can only be arrived at through practical judgment of the particulars of the circumstance. Encounter with instances of a species of value is a necessary prerequisite to appreciation of its broader category.

[35] Aristotle, *NE* 1095a1-1095a5.

We often come to appreciate the value of something or someone only through and because of our first-hand experience with that thing or person.[36] In these cases, it seems as though the appreciation of that value (and the corresponding reasons for pursuing it) inherently depends on having some firsthand experience. In this type of case, something needs to be brought into the debate that the public at large is not in a position to grasp.[37] On the substantive approach, why each of us pursue philosophical study may *look* largely irrelevant to the concerns of morality. This is because, when considered from outside its practice, what it means to be in the grip of serious philosophical inquiry has little to do with our obligations to others. On the formal approach, the value of studying philosophy may fail to appear at all in the language of reasons open to evaluation by all, regardless of their practical experience. If universality is a formal condition for counting as a genuine reason, but appreciating the value of philosophy is conditional on some practical experience, then from the perspective of public justification, the best reasons individuals have for studying philosophy, the ones that grip philosophers when we are inspired, can fail to shine.

If the importance of studying philosophy, that is, the reasons that properly move us when we undertake serious inquiry, cannot be fully appreciated by those unfamiliar with the discipline through generic description alone, then neither the substantive nor the formal conditions on reasons outlined above can always be met. In other words, the full value of this type of activity

---

[36] L.A. Paul (2014) calls these *transformative experiences*. Paul distinguishes between two types: the first she calls *epistemically* transformative experience, whereby an experience gives one privileged access to phenomenological information (about "what it's like") that cannot be had otherwise. The second she calls *personally* transformative experience, where an experience changes how one views or understands the world. The cases I highlight here could be described as either or both types, though it makes no difference to my argument. Where Paul focuses on the potential for rational decision-making about transformative experiences, my focus is on the adequacy of the concept of reasons that sits behind traditional ideas about practical reasoning in moral philosophy.

[37] Interestingly, Hare came very close to this observation: "if pressed to justify a decision completely, we have to give a complete specification of the way of life of which it is a part. This complete specification is impossible in practice to give [...]." (1963, 89-90) Hare concludes that such specification is, in fact, not needed.

slips through the fingers of those outside its practice because fully understanding its relevant value requires first engaging with the discipline oneself. What reason there is to study philosophy may only appear *long after* getting underway.[38] Whether its study has some moral import is to be *discovered*, and often cannot be readily articulated in language suitable for public deliberative exchange.

The same is true of a wide range of important human endeavors. Consider the following case. In the evening hours of July 27th, 1943, the Royal Air Force and United States Air Force combined (in what was code named *Operation Gomorrah*) to drop 2,326 tons of bombs on the city of Hamburg, Germany. The resulting firestorm killed an estimated 35,000 civilians, wounded another 125,000, and displaced nearly 1 million people by destroying over 16,000 multistory buildings.[39] Among the buildings destroyed was the famed gothic cathedral of St. Nicholas church. It is preserved in ruin today as a testament to the loss and grief caused by *Operation Gomorrah*.

Let us start with what is in many ways the least devastating of these losses. Try and think of this cathedral. What exactly was lost when it was destroyed? I am not sure that it is possible to imagine its value clearly without having seen it. Surely, it is familiar and understandable to feel some vague loss when something valuable is ruined. But that value is less familiar. If an art historian were to explain the loss of St. Nicholas by stating that "a great piece of gothic architecture was destroyed", that generic description would fail to capture the magnitude of the

[38] McNaughton (1988, 58) makes this point clearly with the appreciation of jazz. "Once you have begun to understand jazz you hear it in a quite different way. You now have the conceptual resources to detect complex patterns and harmonies within what was formerly merely an unstructured welter of sound. Training has improved your recognitional capacities in music, your ability to discriminate what was there in the music all along. With that increased capacity goes an increasing appreciation of what you hear, an awareness of what is of value in the music."
[39] Apel (2011), 1021. My choice of example is not meant to minimize the atrocities committed by the Nazis in similar bombing runs. See, for instance, their attack on London in the fall of 1940 through the spring of 1941, which resulted in at least 20,000 deaths.

loss. They might do better to say something like "a good example of the gothic style was lost, a priceless work that powerfully expresses the humility and possibility of human worship, a fine example of how space, line and scale can combine to move the human spirit." But to someone who knows little of this place, such generic descriptions can only ever serve as meager approximations of the depth and immediacy of grief that befalls those who have appreciated it for themselves. Something ineliminably tied to direct, first-personal appreciation is lost in the move from the direct experience of individuals to a generic description of aesthetic value. St. Nicholas was a communal place of worship and healing. What was lost was, in part, an understanding of the place this cathedral had in the lives of those dead; its destruction testifies to the destruction of those people, its preserved ruin opens symbolically to the eradication of a whole community. The impossibility of replicating or sharing that particular community, now and forever, begins to uncover why people grieve for this place and why its value escapes formulation in generic terms.

However we hint at it, the value of a particular place cannot be readily formulated in this language because its value lies precisely in the unfolding process of discovery that can only occur through first-personal engagement. To speak in general terms (those terms required to formulate reasons) is to miss the significance of this place for the individuals who lived in and alongside it.[40] As a result, any philosophical model of moral deliberation that aims to take in this value in the form of reasons hits up against this methodological boundary. Like philosophy, the

[40] It is an interesting consequence of my argument that there may be no one left alive who fully appreciates the loss of St. Nicholas. While this may seem curious at first, I think that this consequence fits well with a general idea of how appreciations of a historical tragedy can be dulled by the passage of time.

value of a particular place or work of art is only made available by entering into its practice.[41] Such goods are not to be *given* but rather *unfolded*——doing so necessarily requires effort and time. I submit that this example represents another type of value that features prominently in the deliberative life of many individuals, yet cannot find expression in the model of practical reasoning that relies solely on the articulation of generic reasons for action.

Thus far, we have taken up a handful of seemingly disparate examples. In these cases, it appears that the full value of an activity, experience or relationship escapes meaningful representation in terms of generic reasons. We saw that the value of long distance running, the study of philosophy, place, architecture and even individual loving relationships resists formulation in terms that any rational agent can be expected to find intelligible. We must now ask: what kind of pursuits are these? How wide is this category and what, if anything, do its instances share?

In the cases we have been considering, what completes our understanding of the value of the activity, experience or loving relationship is *doing* the thing oneself. This is reminiscent of what Aristotle calls *energeia*. An energeia is a kind of human activity (also sometimes translated as *actualization*) that actualizes our specifically human capacity for self-directed movement in accord with right reason. Etymologically, energeia derives from the root *ergon* (work), which characterizes the typical function of a (functionally organized) kind.[42] In human life, an energeia

[41] To get a sense for what it means to enter into a practice with a place or a work of art, consider Chi T. Nguyen's (2019) fascinating discussion on how groups can think with and through monuments. Nguyen's view is that art can serve as the vessel for group emotions by making the emotional content of a community sufficiently public so as to be the object of a group commitment. Art, on this view, enables groups to guide themselves with values too subtle to be codified and as such generates commitments for those who come before it.

[42] Aristotle tells us that "the work [ergon] is a completion, and the actuality is the work, hence even the name, 'en-ergeia,' is said with respect to the ergon, and aims at the completedness." Aristotle, *Met.*1050a21–2.

is an activity that embodies the full actualization of the characteristic human function: to direct one's practical thought in a way that expresses a stable and complete conception of the good.

In the examples we have considered so far, what brings us to a complete understanding of the relevant value in question is first-personal engagement. Put another way, *doing* is the actualization of our defining potentialities in that area of life. An energeia is valuable both as the actualization of the human capacity for self-directed activity and as a form of total presence in one's own activity without which they would be dull or robotic. Talbot Brewer is instructive on this point: On his reading of Aristotle, the term energeia

> brings into view the notions of activity as opposed to productivity, and of the actualization as opposed to the withering or malformation of one's most essential capacities. It thereby highlights the fact that for human beings, these two notions are joined at the hip: one's most essential capacities are actualized in the life of complete activity. These capacities remain withered if one is wholly passive, and misused if one devotes oneself exclusively to the production of not-yet-attained states of affairs. For humans, fully actualizing one's defining potentiality—that is, fully becoming who one essentially is—requires the sort of complete and unimpeded activity whose value lies in each of its unfolding moments rather than solely in its expected product. (Brewer 2009, 149-150)

The examples we have considered can be seen as instances of energeia because, when done well, they are more than just ways of passing the time—each can all be deeply held, sometimes lifelong pursuits, whose completion (in the perfectionist sense) can centrally feature in a life

well-lived. I have tried to build up the idea of energeia through consideration of some of its instances. Seeing the examples above as instances of energeia helps to partly explain why their value escapes fair representation in terms of generic reasons. Full appreciation of the value of an energeia requires far more than articulating an intelligible proposition. It requires bringing elements of one's own past and future activities into view simultaneously, to have a continuous, running understanding of the value of that narrative whole. This sort of awareness unifies the self across time by pulling its past and future into its evolving understanding of the nature and value of its evolving activities.

III

However we classify them, one might still insist that the value of these activities *are* accessible to the uninitiated. While it may not be immediately apparent to the novice, the study of place, philosophy or the beauty of nature can, with a little determination, be shared. After all, we must be pointing at *something* intelligible when we speak of art or philosophy as valuable in general, otherwise any example would be unintelligible. The thought is that the problem of capturing these particular values is not a conceptual problem with the role of reasons in moral philosophy, but rather a contingent epistemic obstacle in the experience of aesthetic or philosophical value.

There is something to this objection. Why *can* we speak of the value of these activities in general but not in particular? Answering this question requires a distinction between two kinds of relations between a general category and its instances: these are sometimes referred to as accidental and categorial species.[43]

[43] This terminology is helpfully deployed by Ford (2011), 82-85.

The term *accidental species* refers to relation between a genus and its species where the latter is not logically related to the former. For instance, every broken nose is a nose, but not every nose is broken. What relates the species *broken nose* to the genus *nose* is whatever it is to be broken. Or, put another way, the solution to an equation *nose* + *x* = *broken nose*. The point is that the relation between the genus and species is, in some sense, accidental. We can come to understand *broken nose* by simply adding in some extra feature, in this case *broken*, whose meaning is not essentially tied to the category *nose*. Understanding an accidental species requires merely naming a more inclusive category that is available for understanding by all interlocutors, regardless of their experience. In this sense, the genus is more fundamental than the accidental species. We first explain (or assume) the genus *nose*, and then distinguish the species *broken nose* from everything else within that genus.

In contrast, *categorial species* names a relation where the direction is reversed: we first explain (or assume) the relevant species, and then distinguish the genus from commonalities across species. The paper is *colored* because it is *red* or yellow. But it is not *red* simply because it is *colored*. Our grip on the genus *color*, in this case, is only through our grip on what is shared between its relevant species. Unlike *broken nose*, it is not possible to move down from genus to categorical species by simply adding in some accidentally related feature, in this case we might call it *redness*. There is no solution to the equation *color* + *x* = *redness* because the phenomenal quality *red* is not what it is in virtue of some quality or feature being conjoined with its genus color. Where accidental species describes a relation where some independent quality just so happens to be conjoined with another to create a species within a genus, categorial generality names a wholly different relation whereby appreciation of a genus relies upon familiarity with its species.

This means that we come to understand accidental and categorial species differently. We can understand accidental species through familiarity with its genus, combined with the criteria for delineating its various species. In contrast, the only route to the apprehension of categorial species is via experience of its individual instances.[44] In other words, our grip on this kind of species is *essentially experiential*. In our example, apprehension of *redness* is essentially tied to experience of its instances. Rather than moving downward from the genus by attaching some adjective to the genus and then deriving the species, categorial generality is a relation whereby our understanding moves *upward* from experience with individual instances. From our practical experience, we develop an understanding of the commonality between instances to form a species. From there, we note commonalities between species to form the genus.

I think that something like accidental species is assumed when we object that aesthetic or philosophical value can be spoken of in particular because we can speak of it in general. Given some understanding of the genus *value*, the assumption is that we simply add on some extra feature *aesthetic* or *philosophical*. Capturing *philosophical value*, then, involves only naming whatever is left over after subtracting that general sense of value; that is, whatever it means to be philosophical.

But this cannot be right. As we have seen, coming to understand the value of philosophy or art involves more than simply naming some species of the genus *value* that is accidentally related (i.e. logically independent) of its genus *value*. Instead, appreciating their value looks more like the case of redness, where appreciation of it first requires direct experience of its instances. Just as there is no quality *redness* which, when conjoined with the quality *color* suffice

[44] Christine Tapolet (2004) refers to this distinction as *determinate* and *determinable*. While this terminology has a long history, most prominently in Prior (1949) but dating all the way back to Brentano referring to Aristotle (1982), I prefer Ford's distinction because it sidesteps potential confusions with the connotations of "determinability" that are tied to a narrow array of examples, suggesting that differences in determinability here only come in degrees.

to capture what it is to be red, so too there is no logically independent quality *philosophical* or *aesthetic* which, when combined with the general idea *value* can suffice to reveal the nature of philosophical or aesthetic value.

This point is most clear when we consider the irreplaceable value of particular human beings. It is not possible to capture the value of those thousands of loved ones lost in *Operation Gomorrah* in generic form, absent the direct experience of building relationships with those lost. There is no such thing as a good reproduction of an individual human being. Describing a loved one, even in great detail, can never approximate their value, nor that living history that is a loving relationship. Where a faithful reproduction of a cathedral could perhaps approximate experiencing the original, no such replica exists for, say, one's parent or sibling. Moreover, the intimacy of particular personal relationships is not conceptually sharable—absent first-personal experience, the idea of human intimacy is senseless. Particular relationships represent yet another type of experience whose full value resists being brought into the language of general, universal public reasons because their worth lies in the progressive unfolding of sustained first-personal engagement, precisely the kind of experience that is not available to all.[45]

You might reply that this last point is mistaken. Indeed, you may think that the fact that my examples are intelligible at all proves that such values *can* be adequately represented in generic form. Indeed, you might think that generic description is no obstacle to understanding the value of individual human beings. Put another way, we can know everything we need to know here, in the sense that matters for ethics, from generic description: each human is irreplaceable

[45] Benjamin (2007) makes a similar point about the reproduction of art, though to a different end. For him, the *aura* of a work of art (and as he later clarified all objects), its presence in time and space and its unique existence in its current place is something that cannot be reproduced. However, Benjamin's intentions in formulating the idea are anything but straightforward. See Hansen (2008) for a comprehensive discussion of aura in Benjamin's work.

and equally valuable. Nothing that we could discover in that process of deepening our understanding of individual valuing would provide us with a basis for doubting these claims.

But this is the trouble with generic examples: in order to illustrate the force and point of an example, one must stipulate just enough detail to get the point across, otherwise the example is vacuous. To bring forth a thought experiment is to count on the reader already having some access to the categorial value I'm talking about. Thought experiments assume, in their very structure, that the species of value in question can be put on the table for all to see by merely naming some feature that is logically independent of its genus *value*. But as we have seen, this is not always the case. If we think that what is represented by a generic example is equivalent to real, lived experience, we mistake the model for the thing itself.

What is true of the unavoidable genericity of examples is also true of the language of generic reasons: each can only get us as far as a generic representation of the value of goods in lived experience, never all the way there. Instead, the only route to appreciating the distinctive value of these endeavors is via direct experience with their instances. There is no place for this way of valuing in either approaches (substantive or formal) to the standard view of practical thinking in ethics. This suggests once again that the real issue lies deeper than the content of how we delineate reasons.

What people are actually reasoning *toward* in these cases is made invisible by its representation in theory. On the substantive approach, why we study philosophy, appreciate particular works of art or love particular people cannot be fully rendered in a public deliberative exchange, and must therefore be deemed morally irrelevant. On the formal approach, since the value of such pursuits fails to meet the condition of universality, they may not appear as genuine reasons at all. The issue, it appears, lies with the ill-suitedness of the techniques of philosophical

argument—techniques essentially grounded in an understanding of practical thought that is limited to the articulation and exchange of reasons—for bringing into view some experiences, activities and relationships that properly *move* us in day-to-day life.

The formal and substantive approaches to the standard view assume that readers already have a grip on the relevant categories of value that are outlined by generic reasons (e.g., losing a loved one in general, jazz in general). But such an approach fails to note that our grip on certain values is indefinitely enrichable. When moral philosophy excises discussion of instances from the scene of shared moral reflection, they also cut off a path to a deeper understanding to these species from the business of deeper ethical reflection. As theorists, we deny our ability to deepen our understanding of the urgency of ethics. Without the encounters necessary to get a firm grip on the categories, we lose a grip on the urgency and importance of doing ethics at all. How so?

By classifying the activities whose value escapes fair representation in terms of reasons as instances of energeia, we gained some sense of the expansiveness and urgency of understanding this part of human life. By relying heavily on generic reasons for action for capturing what is worth going for in an energeia, the point and purpose of a whole swath of human life is pervasively left out of the dominant picture of practical thought. Absent a fair representation of the value of energeia, what really enriches our lives in these cases cannot be fully articulated within the bounds of justificatory exchange insisted upon by philosophers. This means that philosophical theories that insist upon the sufficiency of generic reasons will also have to insist upon setting aside an understanding of what really matters in central aspects of actual human lives.

Something has gone wrong. Ethics, if nothing else, is a subject devoted to understanding the materials and resources with which to live a full human life. What could be more urgent than

being able to deepen our understanding of what makes life meaningful? For philosophy to absolve itself from any interest in deepening our understanding of what is valuable in human life, something that each of us is tasked with in our everyday life, is highly questionable. Specifically, it highlights a conservatism at the heart of emphasizing the generic relation: it looks odd that people (philosophers) who devote their lives to the study of how best to live together would sidestep the importance of gaining a deeper understanding of that very thing. And yet, philosophy's distinctive way of representing what matters to actual individuals fails to fairly represent the value of all of their deepest commitments in their full force and form because appreciation of certain goods requires experience with their instances.

How can it be that reflection on first-personal practical experience is, as Aristotle suggests, essential to getting the starting points for ethics, but is also to be exiled from theoretical philosophical reflection? If our theories cannot give sufficient hearing to why we study philosophy, create art or love each other, this is a serious shortcoming. The very good that ethics aims to place at the center of the field (human beings) is also the good whose full appreciation is shunted aside as unnecessary by this methodology, which takes it for granted that a poor representation of that value will suffice. In short, this way of doing ethics *names* the importance of art, philosophy and human beings but does not provide the conceptual tools to bring them into view.

IV

At this stage, it might reasonably be asked: *why* is it so difficult to bring goods like these into the language of public deliberative exchange? My first response was that the dominant versions of the standard view of practical thinking manifest a problematic commitment to

genericity through an exclusive reliance on the language of reasons, which fails to make sense of the experience of particular valuing. I called these activities energeia. Drilling down, I argued that two popular versions of the standard view fail in this way because they mistake the relation between the particular value and its general category—the standard view mistakes accidental for categorial species of value.

There is more to be said here. Perhaps reflection about the value of energeia prompts one to do something that is essentially not generic, public or universalizable; namely, deliberate upon a definite, personal, historically situated life of one's own and the particular people or activities that constitute it. Take the example of love. When I say that I love someone, I am not citing the generic relation as the source of the value I see in them, but rather what is seen from within that relation, the reality of those individual friends whose existence I value. Being a good friend, one might think, requires responding to this person, attending to their needs, not heeding a set of generic principles that stem only from the generic friend relation. To see loved ones as merely *the other party* in a generic relation is to accept their substitutability.

Kyla Ebels-Duggan makes a similar point: "in a large and important subset of cases you could not have fully appreciated that value absent some direct experience. In these cases you could not have come to value the thing as you do merely by considering a report of the reasons or arguments that purport to justify your attitude."[46] To illustrate, she considers the case of love.

> One normally does not, and arguably could not, come to love someone merely by considering claims ascribing admirable qualities to him, or becoming convinced by an argument for the conclusion that one has sufficient reason to love him. Rather, one

[46] Ebels-Duggan (2019), 623. Italics original.

> normally comes to love someone through experience of him, experience in which one appreciates directly the value that one affirms in loving him, that is his value, or perhaps better: him. (Ebels-Duggan 2019, 625)

The point here is relevant to our discussion: the problem with reliance on generic reasons for modeling practical thought is not just that such reasons are imperfect representations of the value of the individuals they purport to represent. The real problem is that representation in terms of reasons, the practice of leveling substantive or formal conditions on their status as legitimate, clears the field of all goods that, despite their prominence in everyday life, cannot be articulated publicly.[47]

I think that Ebels-Duggan puts her finger on an important and overlooked issue here: moral life involves a fundamental fragmentation between different ways of valuing. Such fragmentation is not recognized by unified theories of practical thought. This fragmentation of value represents the development of different points of view (rights, obligations, personal commitments, etc.) that may all be fundamentally incommensurable.[48] Seen in this light, all stripes of moral philosophy represent different attempts to take fundamentally fragmented forms of valuing in human life and unify them by either lumping all modes of valuation together under

[47] Timothy Cleveland (2022, 11) argues that literature can express "things we can know or have insight into that are literally unsayable in a theoretical discourse."

[48] Nagel points to a similar problem: "a person does not feel bound to keep his promises or look after his children because it would be a good thing, impersonally considered. There certainly are things we do for such reasons, but in the motive behind obligations a more personal outlook is essential. It is your own relation to the other person or institution that moves you, not a detached concern for what would be best overall." (1979, 132) He points out that there are formal differences among types of reason and that this difference reflects differences of fundamental types of value, their sources. This, he concludes, rules out a unified solution to conflicts between competing reasons.

the condition of publicity, or excluding some forms of valuing as irrelevant to the domain of morality altogether.[49]

The full implications of this problem have not been recognized. The example above brings out an important difference between the way we appreciate the value of objects and activities on the one hand and persons on the other. Ebels-Duggan's insight hints, but does not explicitly say, that there is something singular, and unreproducible about the way we value individual people. As noted above, where first-personal experience of activities and objects are often roughly replicable (i.e., others can share the same object or begin the same activity type) there is no corollary process of reproduction available in the case of friendship formation or other loving relationships. In the case of human relationships, this is partly because people are always works in progress. A central feature of human life is that we grow and change through time. To be in a loving relationship is to put oneself in a position to witness and participate in that growth, which requires a different dynamic relationship than we have to static objects. This is partly why people are irreplaceable. Where a cathedral can be rebuilt, returned to again and again as a fixed point of reference, human beings are each *sui generis*, dynamic and constantly unfolding. As such, their value cannot be fully presented in the form of a consideration whose meaning can be made available to all just by virtue of comprehending a general description.[50]

---

[49] Nguyen (2022) provides an interesting analysis of what happens when this one-size-fits-all approach becomes internalized by the agent themself: he calls it *value capture*. Value capture occurs when our natural values, which are rich, subtle, and hard-to-express are displaced by social or institutional pressures, which present simplified, typically quantified, versions of our values back to ourselves. Our values are said to be "captured" when those thin, commodified versions of our values become part of our deliberation. Nguyen notes the familiar example of starting to exercise for the sake of your health, then getting captured by a fitness tracker and gradually coming to care only about your daily step-counts.

[50] A similar point is made by Anderson (2017). Where the reasons offered in moral justification are limited to finite definite descriptions of general relations, the reasons of individuals are inexhaustibly describable because they stem from specific persons. This echoes Murdoch's (1997) conception of the infinite task of appreciating the reality of individuals. See also Hopwood (2017) for a thoughtful elaboration on Murdoch's work on moral concepts.

For her part, Ebels-Duggan seems optimistic that such fragmentation in the way we value could, if properly modified, be incorporated into Kantian form. She suggests that a reinterpretation of Kant's moral psychology to include an expanded picture of inclination might do the trick. For Kant, inclinations call out for causal explanations and not normative ones. After all, an inclination is a spontaneous attraction, and hence does not appear to be something *we do*. Kant thus divides all motivational states into one of two categories: those attitudes for which we are able to give some reason on the one hand (e.g., moral motivation), and those for which we could have no reason on the other (e.g., inclination). But as Ebels-Duggan points out, love does not fit neatly into either. For, it seems wrong to hold that we have articulable reasons to love a particular person, but equally wrong to say that such affection is random or without normative commitment. She then suggests an expansion of the class of motivational states beyond those strictly endorsed by Kant.

So conceived, this new understanding of Kantian incentives would expand the class of inclinations to include things that agents do, and so—like actions and unlike headaches—things that they can aptly be asked, or ask themselves, why they do. However, neither those inquiring nor the agents who are the targets of this inquiry should expect that a full answer to this question can be given in propositional form. Unlike the incentive of pure practical reason, these would be incentives we cannot reason our way into.[51] Her idea is to expand the Kantian understanding of motivation by inclination beyond arational, passive states to include the appreciation of values for which we cannot articulate reasons but which nonetheless reflect the active endorsement of normative commitments. If a Kant-inspired theorist could give an account of inclination that

[51] Ebels-Duggan (2019), 638.

could include these features of direct appreciation, then we might be able to vindicate the Kantian model of moral deliberation.[52]

I am not sure that I fully understand the sense in which this view is recognizably Kantian. This is because Kant's account of moral deliberation leaves little room for inarticulacy about reasons. In deliberating about what to do, the Kantian moral agent seeks to act only on maxims that could be willed universally. To see whether this is possible is to investigate whether one's commitment can be intelligibly challenged or pressed for justification. A fully justified agent is one who is capable of answering that challenge in the form of identification of a rational condition for the commitment in question. To fail to affirm the condition in question is to fail to rationally maintain confidence in the original commitment. If, upon further questioning, the rational condition is itself subject to further legitimate challenge, then the need for justification repeats. As a result, the only way to be fully rationally secure in one's commitments, on the Kantian view, is to trace such a line of reasoning back to commitments that are not subject to any real challenge.

But to go this far is to accept a version of the formal condition I have brought into question. Recall that the Kantian process of rational justification is predicated on the condition of universalizability. Such a condition is difficult to parse absent a foundational commitment to universality and publicity. The ability to articulate one's grounds for action is built into the very notion of a categorical imperative. In Kant's view, absent the ability to articulate the considerations that ground one's action in publicly affirmable form, the moral agent is simply not justified. This makes it very hard to see how to assess the claim that some exception is needed to

[52] Nauckhoff (2003) develops a close corollary to this view, arguing that respect for the moral law is a "non-pathological feeling" that arises necessarily upon reflection on the moral law. However, she does not go so far as to expand Kant's notion of inclinations in the way Ebels-Duggan imagines.

an otherwise binding duty because of this aspect of this particular activity. The weight of the case for the exception cannot be put on the table. If so, it is difficult to see how we can have confidence in a verdict that has been rendered without an understanding of its ground.

This leads us to a surprising conclusion. If we accept Ebels-Duggan's observation about direct appreciation and its implications for our ability to articulate reasons, as I do, then there is little hope for reconciliation with mainstream moral theorizing (not only Kantian forms) because the full articulation of reasons in the form of a (substantive) publicity or (formal) universality condition is among its foundational commitments. I think that this issue is worthy of more attention than philosophers typically give it.

The issue I am highlighting here would be perplexing for Kant. On his view, what we must recognize, in order to get on with the business of doing moral philosophy, is that the rational nature of human beings is worthy of the utmost respect, and that in this respect, each of us is *equally* deserving. Respect for the rational will, combined with equality before the moral law, is enough to get moral principles off and running. A Kant-inspired moral philosopher, then, feels the philosophical pressure to make our rational nature the ultimate object of respect, but has not yet revealed (or attempted to reveal) the respect-worthiness of the rational will itself.

To be clear, I am not denying this assumption (that rationality is worthy of respect), but rather highlighting the fact that this brand of moral philosophy is, by and large, not interested in bringing into view the respect worthiness of rationality itself. Similarly, Kantian philosophers may be quick to agree, under pressure, that those lost to the tragedy of war *are* irreplaceably valuable, but at the same time show very little interest in deepening their understanding of *what* is so irreplaceably valuable in human life. To account for the good of an activity, experience or person, reasons must make available for all both a conceptual grasp of its nature (character and

meaning) and a contextual understanding of its significance (depth and status in one's life). I have tried to show that in a centrally important range of cases both elements, the nature and the significance of particular human goods are either distorted or left out by the methodological constraints built into the dominant views of generic reasons.

The result is troubling. If appreciating the value of some of our deep normative commitments requires a unique, hands-on experience with the object, activity or person in question, then it follows that a mere report of what one found worth going for in the activity or thing or person will be insufficient. Articulating a generic version, one fit to fulfill a substantive restriction on what counts as morally relevant or a formal condition on reasons more generally, is more than just to give an incomplete report; it is to sweep away any and all of considerations that feature prominently in personal moral deliberation *simply because they cannot be shared*. To do so because such considerations cannot be given a hearing in language that all can understand, to exclude them from the domain of legitimate objects of practical thought, looks arbitrary. If moral philosophy earnestly seeks to make substantive claims about what humans ought to do, it cannot ignore or distort such familiar and important human pursuits.

While some have recognized this problem, none have addressed its radical implications: attempts to produce tidy theoretical accounts of deliberation and practical thinking will systematically fail to fairly represent many of the experiences, activities and relationships that we care deeply about.

V

Suppose one were to grant the argument so far: that there *are* goods in life whose value cannot be articulated in the necessary form in order to meet the conditions on reasons highlighted above.

One might agree that there are ways of valuing that resist being brought into the language of public deliberative exchange as reasons for action, but still argue that this makes no practical difference in normative ethics.

Why not? The proponent of this view might argue that the exclusion of certain goods from representation in terms of reasons in public justification for action does not entail that they are not legitimate in other contexts. The substantive approach we have been considering merely incorporates, into the framework of practical thought, a practical distinction that enables us to deal with the inevitable limitations of our moral capacities and the complexity of our social circumstances. We might acknowledge, so the argument goes, that there are goods in human life that do properly function as reasons in our personal lives, but that cannot be sufficiently articulated in the language of *public* deliberative exchange. As a result, we can rightly exclude these personal considerations from the domain of those considerations fit for theorizing about what is right and wrong for all. We can do this because there are real, practical social problems to solve, and so we must restrict the domain of relevant reasons to only those where we can rightly hope to find common ground.

The appeal of this response is linked to a conception of moral deliberation as a societal commons whereby the theorist's ambition is to ratify only those considerations that can be given in a public deliberative exchange. This last point is significant. The scope and scale of the public involved in moral life—together with the depth and seriousness of their differences—makes outlining the reasons we have for action (and acceptable standards for what can serve as a justifying reason in that context) both daunting and urgent. Rawls makes this clear at the outset of *Political Liberalism*:

> A modern democratic society is characterized, not simply by a pluralism of comprehensive religious, philosophical, and moral doctrines, but by a pluralism of incompatible, yet reasonable, comprehensive doctrines. […] a society may also contain unreasonable and irrational, and even mad, comprehensive doctrines. In their case, the problem is to contain them, so that they do not undermine the unity and justice of society. (Rawls 1993, xxvii)

Rawls takes the fact of reasonable (and unreasonable) pluralism about the human good to be a serious threat to any practical system of ethics. How, one might ask, are we to live together without limiting, censoring, harming or even killing one another given the stark reality of the world? After all, most conflicting, comprehensive systems of ethics do not readily admit of compromise.[53] History books overflow with examples of our collective failure to do so: the Reformation, the Spanish Inquisition, the Holocaust are just a few that Rawls notes in passing. The implication is that while it may appear arbitrary at first glance to rule out certain modes of valuing as irrelevant to morality simply because they cannot be shared, Rawls' political philosophy suggests that we have overwhelming reason to do so. Where I have insisted on the importance of acknowledging values that cannot be readily shared in the language of reasons, a Rawls-inspired account of the point and purpose of ethics takes its starting point to be just the opposite: to set substantive standards for what counts as moral justification, to exclude certain types of personal values like religion, personal affection for individuals or other idiosyncratic attachments, precisely because they cannot be shared.

[53] Rawls (1993), xxvi.

What might such a view look like? While Rawls does not offer a comprehensive moral theory, we can get a glimpse by examining his limited work on Kantian constructivism. For him, "the task of a moral doctrine is to specify an appropriate conception of the person that general facts about human nature and society allow."[54] Central to this conception is that the only considerations that are relevant to the project of creating a moral theory are those that are *public*. The idea of publicity has multiple levels in Rawls' political theory,[55] but for our purposes we can think of publicity in this sense as a constructed social point of view that is authoritative with respect to all individual points of view.[56] That point of view acknowledges the free and equal status of individuals, in part, by only making considerations binding if they can be acknowledged by all in virtue of some shared background ideals. "Whether certain facts are to be recognized as reasons of right and justice, or how much they are to count, can be ascertained only from within the constructive procedure, that is, from the undertakings of rational agents of construction when suitably represented as free and equal moral persons."[57] He tells us that it is up to those contracting parties to decide how simple or complex the moral facts are to be, that is, to decide on the number and complexity of the principles that identify which facts are to be accepted as relevant reasons.[58] This enables "all members of society to make mutually acceptable to one another their shared institutions and basic arrangements, by citing what are publicly recognized as sufficient reasons, as identified by that conception."[59] Just as the fundamental terms of social

---

[54] Rawls (1980), 534.

[55] Rawls outlines three levels of publicity. The first marks a recognition by everyone that society is regulated by the principle of justice that is known by all, a fact which is itself also known by all. The second marks a set of beliefs about institutions and human nature that flow from these agreed upon principles of justice that are also accepted by everyone such that their wide acceptance is also public knowledge. The third marks a complete justification of a public conception of justice such that anyone in a given society, if motivated to reflect thoroughly enough, could articulate its form by pointing to their shared culture, laws and political institutions. See Rawls (1980, 536-537).

[56] *Ibid*., 554.

[57] *Ibid*., 519.

[58] *Ibid*., 565.

[59] *Ibid*., 517.

cooperation between free and equal moral persons should answer to the requirements of publicity, so, too, must the fundamental terms of any adequate moral theory. In order to hold each other accountable for our actions, it might be said, we must be able to share the grounds of our choices. Reasons must be public in this sense because the *question of how to live together is public*.[60]

Given the fact of reasonable and irreconcilable pluralism about the good, a moral theorist might begin instead with the right; such a view takes our overlapping concerns for safety, liberty and the provision of what Rawls calls "primary goods" as starting points for determining the substantive limits of our moral duties. Seen in this way, publicity becomes a substantive condition on what counts as a normative reason: whether a consideration can be shared by others is what *makes* it morally relevant—if a consideration is by its very nature only relevant to a single individual, then it falls outside of the domain of practical moral life.

This is a powerful idea. But even so, I do not think this type of response works as a response to my worry. I think that the conception of moral theory offered above is fundamentally inadequate because it bars reflection on the phenomena of individual valuing by contracting the scope of moral concerns that can be stated generally; it deems *insight* into those general concerns, which fill in the content of the proper relations with our fellow human beings as irrelevant, even compromising for theories of that very topic. Such insight, I contend, is crucial to full engagement in human life.

To see why, recall my initial example—choosing spousal care over tax forms. On a Rawls-inspired view, it is unclear whether morality ought to have anything to say about this case.

[60] Rawls is quite explicit: "In public questions, ways of reasoning and rules of evidence for reaching true general beliefs that help settle whether institutions are just should be of a kind that everyone can recognize." (1980, 539)

If morality is about setting the principled terms for social cooperation and those terms must be set in language that can be shared by all, then it seems that neither my obligations to the government, nor my wife can get more than a cursory hearing. If anything, a Rawls-inspired constructivist account of morality might say, in general, we have good reason to fulfill our duties as citizens (like paying taxes) in order to maintain the smooth administration of social services. But short of upholding the general, fundamental terms of social cooperation, which apply to all agents alike (e.g., prohibitions on harm, threats of violence, theft, etc.) such an account is silent on the value I see in my relationship. The unique details of individual personal relationships are simply outside of morality's purview.

On the Rawls-style approach, this is just as it should be. After all, it seems reasonable to hold morality ought not to invade our personal lives. On this view, since the obligations that flow from the fundamental terms of social cooperation have no place in our personal relationships—again, so long as we uphold and respect the terms of social cooperation that apply equally to all—there is no practical need to represent those values in a moral theory. In the realm of the personal, we can (and should) do as we please.[61]

Why, then, do I insist that a moral theory *ought* to provide the conceptual tools to fairly represent the values I apprehend through my wife and our shared relationship? Consider again what it would mean to *deny* that such values are morally relevant. At bottom, the practical, substantive conclusions of a moral theory purport to have normative force. If so, then my ignorance, transgression or defiance of these conclusions is, at least, a failure on my part to take

[61] My criticism here parallels an important feminist critique of Rawls. See Held (1987); Nussbaum (2003); Schwartzman (2006). Because social contract theory has a history of leaving the questions of familial labor, rights and justice unanswered by relegating them to the private sphere, Rawls' defenders (and even Rawls himself) have conceded that the problems of "justice of the family, the equal justice of women and how these things are to be achieved" (Nussbaum 2003, p. 488) are particularly difficult for liberalism.

seriously the claims that others have on me. But the reverse can also be true. I find Raimond Gaita's writing instructive here: for an agent to justify himself to others "[s]omeone must not only take seriously what he claims to be of moral concern to him; his taking it seriously must be capable of being taken seriously by others." To make this possible, Gaita concludes that a "morality must deepen rather than cheapen our understanding of what we care for." This is because, we cannot say "that it is irrelevant whether morality deepens or cheapens our understanding of our lives without making it unintelligible what morality's claim on us ought to be."[62] What is the cheapening here? We cannot settle questions about the contours of our duties unless we are operating with some picture of the actual values that are in play. Over time, the effect of attempting to avoid these deep questions is a shared public representation of our grappling with those deep questions that cheapens rather than deepens. To suppose that this is not so is to suppose that understanding, inquiry and meaning making is somehow outside the bounds of morality. Following Gaita, and contrary to a Rawls-style substantive approach, in order to have an intelligible claim upon what we rightly do, a moral theory must at least intelligibly claim to deepen rather than cheapen our understanding of what human life entails and provide language with which we can make ourselves—in all our subtle, beautiful and serious complexity—intelligible to each other. And to do *this* work, it must offer a deepened perspective on the facts of human life, the good and evil we might actually suffer. To use Gaita's phrase: it is to hold that morality is one thing and the meaning of things another entirely.

But surely drilling down to the meaning of actual human relationships, philosophical study, our relation to art and place *is* an integral part of morality because they are integral parts of our lives. And if the conclusions of a moral theory are to hold sway with me, to instruct what I

[62] Gaita (2004), 38.

ought to do, then it must be possible to use its language to both gain insight into what I care about, and yield some sense of the seriousness of that insight for others. The language of practical reason that stands behind the standard view does not yield such a possibility.

Let us ask once more: what does it mean to deny that moral theory ought to capture the value I see in my personal relationship to my wife? It is to also deny that the value I see ought to concern other people, that they ought to take the considerations that flow from it seriously. To affirm this is not to argue that the moral theorist must take personal relationships as overriding, nor that they must first recognize the concerns of every individual on the planet before moving forward. My point is only that accepting the terms set out by a Rawls-style approach, the terms of what moral theory must be and do is to accept a view that cheapens, not deepens our understanding of what is worth going for in human life.

I have suggested that there is a way of valuing certain familiar and centrally important experiences, activities and relationships, that resists formulation in the form required to meet any publicity or universality condition on generic reasons. Because generic reasons are taken to be sufficient to capture what is significant in practical decision making, such goods as the ones I have outlined are given no place in the dominant approaches to modeling practical thinking, despite their central place in the lives of real people. The problem is that this is not a trivial class, easily cast aside as irrelevant to life. Indeed, if we cannot even put on the table the very value that lies at the heart of moral theorizing, the irreplaceable value of human life, this calls for a reevaluation of our philosophical methodology.

You may still have doubts. And it is not easy to address these doubts directly. To combat this, I implore the reader to look to your own life: hold in your mind an object, activity, experience or relationship of immense value to you. Now attempt to articulate that value, fully

and clearly in words that form a reason to act on its (their) behalf. The good in question, seen from the inside, is much more urgent and essentially tied to the particular historical and social circumstances of your life. All this information, which can only be accessed via the progressive unfolding that marks an attempt to engage with this activity or experience of value, is left out of its formulation. My argument attempts to center those details that are only observable to you. Indeed, you may be able to know more than you can express to others, and some of what you know to be worth caring about will be beyond your capacity to bring into words entirely. It would be a misstep for me to offer any philosophical thought experiment that attempts to bring all my readers around. I think it is a consequence of my argument that its vindication must happen by the reader's active attempt to bring its details into view.

VI

Where does that leave us? My aim has been to highlight what I see as a serious oversight at the heart of contemporary moral theorizing. In particular, there exists a class of familiar and important human values that resist being brought under the language of public deliberative exchange that characterizes theoretical attempts to model practical thought. What falls under this description are not esoteric whims, but deeply familiar human activities: appreciating art, beauty in nature, philosophy or particular human beings are just some examples of activities whose appreciation is irretrievably tied to first-personal experience. Since these are allegedly *central* values in the study of ethics, it is cause for concern that their full appreciation is either ignored or downplayed by the methodology we have considered. Talk of generic reasons for action provides philosophical language that deflects the difficulty of bringing the value of these pursuits into view. Since the goods inherent in these pursuits are only available for true appreciation after

undergoing a first-hand attempt to undertake them, they are not fairly represented (if shown at all) in philosophical models of practical thinking that take publicity or universality to be a fundamental attribute of reasons. That this problem of exclusion is common to both dominant approaches to representing reasons in ethics suggests that the heart of the issue is not the *content* of these approaches, but rather the philosophical methodology that relies heavily on generic reasons to model practical thought.

By way of conclusion, I'd like to point out two more problems raised by my claims in this chapter. I have argued that the very *method* by which philosophers attempt to faithfully represent the concerns of individuals includes either a substantive commitment that misses a whole range of important values that properly motivate human beings or a formal condition that passes over these values entirely. If there is such a class of familiar and important experiences, activities and relationships whose value cannot be fairly represented, it follows that any subsequent account of practical deliberation—whose quality depends precisely on the fair representation of all relevant considerations—will be correspondingly stunted. To reject the necessity of appreciating the value of philosophy, art or loving relationships in building moral theories is to restrict the powers of our imagination, to decide to set a final limit to the indefinite development of our moral intelligence.[63] The descriptions of individual valuing that are rejected as irrelevant to philosophy are just those that deepen our moral understanding. Where philosophers, by and large, seem content to say *that* philosophy, art and human beings are valuable, they seem to think that what it means for these goods to be valuable is somehow less important.

Framing the issue this way reveals another problem. If practical thinking can only be represented in the language of genericity, this requires individuals who attempt to represent their

[63] A close account of the development of our powers of imagination will be a primary subject in chapter three.

own experience to take a certain detached stance toward their own personal projects and commitments. Where many believe this to be a welcome feature of the theoretical accounts we have been considering, I think this should be cause for concern. Seen in this way, moral theorizing encourages a certain objectionable way of being in touch with other people, one that guides us to see ourselves and our loved ones through the lens of the impartial moral perspective. This tension between how we value what we hold dear and how we must describe it in order to be justified gives rise to a seemingly intractable tension between the impartial demands of morality, and partial demands of values like familial relations, love and friendship. To be a morality worth articulating, for it to lay claim to prescriptions of what we ought to do, it must (at least plausibly) claim to deepen, not cheapen our understanding of what is worth going for in life.

Further discussion is necessary to probe the full depth of these issues, but I will offer one suggestion here: doing ethics in the way we have outlined above manifests an implicit refusal to acknowledge that ethics cannot be *finished* on paper, an unwillingness to leave anything relevant to the argument incomplete. That aspiration, when taken to its extreme, ends up being an affront against the status of our fellow humans, for it disregards their role as sharers of an essentially personal task that they are also in the midst of. This amounts to a questionable refusal to engage in the enrichment of the categories of value that philosophers devote their life to trying to sort out. My argument calls out for an offer to be made from author to reader—an offer to engage in the selfsame task that ethicists recommend as worthwhile. This is an entirely different spirit of doing philosophy, one that centers the limits of our theoretical ambitions.[64]

[64] Gaita (2004, 13) sheds light on the spirit I have in mind: "I do not wish to prejudge the relation between reflection within the subject and reflection outside of it, although I plead for greater philosophical patience [...] I shall try to undermine the confidence on the part of philosophers that they know what to make of it, that it is for philosophy to delineate all the serious options and that what is said outside of philosophy will, at best, speak for one or another of them."

Chapter 2: Reasons, Individuation and Practical Thinking in Dancy's Particularism

I

In the previous chapter, I outlined a limitation of what I called the standard view of practical thinking, whose central focus is on actions done for reasons. I examined two common approaches delineating reasons within the standard view. We saw that each of these approaches systematically exclude an important class of human experiences, activities and relationships from fair representation. This class is broad, ranging from studying philosophy to appreciating art to being in love. While obviously central to practical thinking, why we pursue these things essentially resists being heard in the language of wide-open debate. By insisting that any genuine episode of practical thinking is a matter of identifying reasons for particular actions, the standard view artificially restricts the range of activities that make up practical thought. I concluded that what is worth going for in philosophy, art or a loving relationship cannot always be given in the form required to meet a substantive condition for moral relevance or a formal condition for normative reasons more generally.

But there is another tempting response to my argument from the previous chapter. Perhaps the problem of unjustly excluding certain goods from the domain of practical thinking stems from a problematic commitment to genericity or universality in philosophical theory building. If so, then why not simply avoid genericity and universality altogether? Maybe, so the response goes, the source of the error boils down to an insensitivity to the uniquely individual experience of certain activities, experiences and relationships. If this is right, then an account of reasons that depends only on the agent's sensitivity to the unique features of a situation is a natural solution. Particularism about reasons represents just such an appealing alternative.

This chapter outlines why particularism about reasons looks to be a natural solution to our worries from chapter one. Focusing on the work of Jonathan Dancy, it then argues that despite this appeal, particularism's most prominent advocate fails to address our methodological worries because his view shares some foundational assumptions about the role and function of reasons with the generalist views we have already discussed. In particular, this chapter argues against the assumption that reasons can always be meaningfully individuated from their background context that provides their normative significance. The chapter concludes by illustrating how the failure of this assumption, and others that flow from it, together have the further consequence of overlooking a compelling objection to generalism. What is revealed is that the debate between Dancy's particularism and generalism about reasons in moral deliberation only begins after another, more fundamental set of methodological assumptions about practical thinking are already in place.

If successful, this chapter will provide two contributions. By uncovering the ways in which Dancy's version of particularism aligns with generalism, and as such fails to address our initial concerns about practical thinking, our investigation will light the way toward a new particularism that can. In doing so, it will also reveal that the structure of practical thought is not exhausted by talk of reasons for action, which will be discussed more in the following chapter. In short, while Dancy's particularism about reasons embraces the uniqueness of individual experience and rejects the universalizing impulse of most analytic moral theorizing, I show why its radical revision of the methodology of modeling moral deliberation *does not go far enough*.

II

Ethical generalism is the view that moral judgements depend on a suitable provision of exceptionless general principles. *Ethical particularism* is a family of views that denies this generalist thesis. Different versions of particularism are defined, in part, by how principles are conceived. As a result, there are at least as many versions of particularism as there are conceptions of principles.[65] And while it may be a useful practice to carefully untangle them all, my ambition lies elsewhere. Rather than a comprehensive review of particularist literature, this section will focus on the work of Jonathan Dancy as an exemplary, if unique, proponent of particularism.

At first glance, it looks as though Dancy's account of reasons could not be more different from the generalist. Where the generalist demands sameness in the way a consideration functions across different cases, the particularist insists on variability. The former view is sometimes referred to as *atomism*. According to atomism, reasons are considerations whose status and polarity remains the same across cases.[66] As the name suggests, atomism assumes that it is always, in principle, possible to evaluate the status and significance of reasons individually, one by one. Thus, atomism is both a conception of how reasons function in theory across cases (static, unchanging polarity) and a way of understanding how we can come to understand and evaluate them within a given case (as individual entities).

---

[65] For instance, there are arguments against principles as *absolute* or *contributory* determinants of right and wrong. (Dancy 1993; 2000; 2004) There are also arguments against seeing principles as *standards* of any kind. Armstrong (2004) holds a version of this latter view, claiming that principles are truth-making standards for moral propositions. Another view holds that principles are *guides* to right action. (Railton 1984) Yet another approach combines these views, arguing that while principles purport to explain the moral rightness or wrongness of actions, they do so in a way that is not particularly well suited to action guidance. Kant seems to hold a version of this view.

[66] See Dancy (2004), 73–4; Dancy (2000), 130; Dancy (1993), 60.

In contrast, Dancy holds that whether and to what extent a consideration counts as a reason is a function of a holistic evaluation of the context in which it occurs. This thesis is referred to as *holism* about reasons, and it forms the core of his particularist moral insight. But despite its importance in his view, Dancy is less than clear about what holism entails. At minimum, holism indicates a commitment to a context sensitivity in moral judgment. But what makes this a uniquely *particularist* commitment? After all, it is uncontroversial that moral judgments depend upon some sensitivity to the demands of a given context. In this sense, moral theorists of all stripes sign up for the context dependence of reasons. For instance, all might agree that the recommendation of a friend may in one case count in favor of trying a new activity, but circumstances could alter or even reverse the direction of favoring. If that friend is suffering from delusion, lack of information or is under the spell of a pyramid scheme, that same recommendation now carries a radically different import. Nothing about this context sensitivity rules out generalism.[67]

Perhaps what separates Dancy's commitment to holism is the *extent* to which context matters in determining the status of a consideration. For him, whether and to what extent a friend's recommendation counts in favor of doing some action can only be determined on the basis of a sensitive evaluation of the features of a given situation as a whole; its status cannot be determined when detached from that context. For Dancy, this does not mean that reasons themselves, on the particularist view, should be understood so expansively as to incorporate, in their content, everything relevant from the whole context. Instead, a plausible interpretation of context sensitivity holds that an adequate understanding of when a consideration counts as a

[67] Dancy would describe this case in terms of *enabling conditions*. For example, an enabling condition for a promise to count as a genuine reason might be that it is uncoerced. To this Dancy rightly adds that it would be futile to attempt to write down every enabling condition for every consideration.

genuine reason always requires one to look toward the larger whole of the context in which the candidate reason appears.[68]

One might still hold, however, that this demanding claim of context sensitivity is not enough to distinguish particularism about reasons from generalism. This is because the generalist might be happy to sign on to a thoroughgoing context dependence, even one that holds the status of reasons to be a function of the holistic consideration of the relevant features of a particular situation.[69] Simply holding that a deeply rooted context sensitivity is necessary for responsible moral judgment, by itself, does not seem to rule out the possibility of general moral principles.[70] Distinguishing a uniquely particularist version of holism, it seems, requires the addition of some other thesis.

Another path might be gleaned from a consideration of how to *codify* holistic evaluation. On one natural interpretation, holism looks like a very complex version of atomism. This reading holds that if a moral agent were to consider a particular case as whole, in all its complexity and layered detail, then it seems unlikely at best that a consideration's meaning in context will be remotely the same across different cases. This is because any given consideration within its own unique situation is dependent for its status as a reason on a *complete* specification of all the relevant details of that case, If we are to take seriously the assumption that the world is complex and varied as a function of radical differences in social circumstance, cultural heritage and

---

[68] Dancy (2004), 73–4; 95-99; Dancy (2000), 130; Dancy (1993), 60.

[69] Scanlon (2014) seems to go this way when he claims that the status of normative facts can only be determined from within a first-order, substantive moral deliberation.

[70] Many have argued that a strong commitment to context sensitivity is all there is to holism about reasons, but draw very different conclusions from this point. For instance, Cullity and Holton (2002) argue for a weak particularism that cites "presumptive reasons", which mark a class of considerations whose normative valency varies depending on the context (whereas some considerations do not). Stratton-Lake (2000) denies that there can be any principled relation between certain natural and moral properties that mark them as invariant reasons (except by historical chance). See also (Little 2000) for a version of this view.

individual psychological history, it seems plausible, so this version of particularist claims, that the context sensitivity of a consideration transcends any useful codifiability.[71]

But once again, this natural reading is misleading. If holism about reasons were simply a complex version of atomism, there would be no philosophical problem to solve, but rather a contingent, anthropological one. Seen in this way, holism amounts to the view that what is required to make adequate moral judgments is simply a *whole lot of data* about the situation at hand. While knowing more about the relevant context is surely useful, this is not, to my mind, the target of a particularist version of holism about reasons.

What, then, is meant by particularists, like Dancy, who deny the codification of reasons? The point seems to be that features of the world do not always ground the same moral meaning. This is because the moral contribution these features make on each occasion is itself dependent, in a way that escapes useful or tractable articulation, on what *other* nonmoral features are present or absent, considered holistically together.[72] This sense of anti-codifiability does not mean that cashing out the propositional context of particular situations is too hard, but rather that *there is no way* to cash out the context in which the moral meaning depends in tractable propositional form.[73] While some have argued that both demanding context sensitivity and the anti-codifiability

[71] Like context sensitivity, anti-codifiability is sometimes taken to be *the whole* of particularism. If it is true that generalists can sign up for the radical context sensitivity that is popular with particularists, then it seems that the only structural feature of holism that provides support for particularism is the anti-codifiability thesis.
[72] See Little (2000), 280.
[73] The form of holism I am discussing here is sometimes referred to as *unrestricted holism*. This is the view that moral considerations are, in some sense, always open to new interpretation on the basis of new information. Jackson, Pettit, and Smith describe this form of holism as the claim that "no matter the quantity and nature of the descriptive information you have that provides a reason for some moral conclusion, say, that X is right, more may come to hand that leaves the previous information undisturbed and yet, when combined with it, provides a reason against X being right." (Jackson, Pettit, and Smith 2000: 99)

thesis are *both* compatible with generalism, for our purposes, we can take these two theses as central pillars of Dancy's particularism about reasons.[74]

Dancy takes what I have called demanding context sensitivity and anti-codifiability to be two theses that suggest a third: absolute principles are at best theoretical crutches—and at worst fictional delusions—which are not needed to model moral deliberation. If reasons vary radically in strength, relevance and polarity across different cases, and do so in ways that cannot be codified, then there is no need for absolute principles to specify any special feature(s) that always succeed in making (or contributing to) an action or attitude being right or wrong—no such features exist.

The three theses above also serve as the basis for his own particularist conception of moral deliberation and public justification. "[t]o justify one's choice is to give the reasons one sees for making it, and to give those reasons is just to lay out how one sees the situation [...]."[75] Where the generalist locates moral justification in the specification of right-making principles, Dancy sees only considerations that are always inextricably linked to features of a particular situation whose normative authority does not derive from general principles of any sort. Where the generalist conceives of the task of moral theory to be the specification of right-making general principles, Dancy identifies the considerations that bind us in moral obligation to be inextricably tied to the features of each unique circumstance. Dancy's particularism, then, is also

[74] For a detailed discussion about the relation between holism and particularism, see McKeever and Ridge (2006), especially chapter 2. Their conclusion is that we can be particularists in our individual practical judgments, but use generalism as a regulative ideal. Since our aim is not a systematic investigation of holism or particularism, we can safely pass over controversies about how to interpret the relation between the two. Instead, our focus is on how the particularist position, as generally conceived, might answer our worries about artificial exclusion of certain modes of valuing from moral deliberation.

[75] Dancy (1993), 113.

a radically different way of seeing what moral life is; one that always essentially involves reckoning with the features of the particular situation at hand.

At this point, we have seen that Dancy's particularism offers what appears to be a promising set of claims about the nature and significance of reasons. He holds that reasons must be evaluated holistically, with a demanding, context-dependent framework whose relations cannot be codified or represented in the abstract. Dancy takes these claims to cast doubt upon the existence of general principles that claim to specify features of the world that fix the status, significance and polarity of reasons. This view subsequently fixes a conception of justification that requires specification of reasons-giving considerations considered as a part of a whole contextual situation, one that is always conceptually subject to modification on the basis of new information. All this marks Dancy's particularism as a potentially powerful response to our worries from chapter one. But before we examine how deep this apparent power goes, we must first investigate how particularism's most prominent proponent might answer our concerns from the previous chapter.

III

How might Dancy make sense of why we study philosophy, make art and seek love? What does Dancy's holism about reasons offer that the atomistic generalist cannot? Holism's commitment to context sensitivity suggests that one's individual reasons might be radically distinct from others. What draws *me* to philosophy is a function of the holistic situation in which I find myself as a newcomer. This situation can be seen as a complex series of situational contingencies: what texts have I encountered? What themes do they take up and in what tenor? With whom do I happen to study? What classroom dynamics exist and how do they encourage or

discourage certain psychological tendencies that I already have? How does all this combine to pique my interest, which could then spur me on to further study? All this and more fill in the features that ground the reason-giving considerations in the example.

Holism's demanding context sensitivity looks like an appealing start toward answering our worries about artificial value exclusion because it does not suggest, on the surface at least, that the reasons that exist to study philosophy are necessarily *the same for all*. Moreover, since the philosophical novice, by definition, does not have mastery over the methods or content of philosophical study, it follows they will likely be unable to attend to the minute details that ground what is important about philosophical study. Dancy might argue that becoming initiated into the practice of philosophical study *is* learning to be sensitive to the very contextual details that make up an encounter with philosophical argument. Absent mastery of this important skill, the reasons for pursuing it will be understandably overlooked.

Holism's commitment to anti-codifiability makes the picture look more appealing, still. Recall the central worry with this case: why one begins to study philosophy may not only be unavailable for articulation to others, it may not even be apparent to the agent attempting to study it. On the standard models we considered in chapter one, this opacity rules the goals of philosophy out of contention as candidate reasons. Anti-codifiability provides a natural explanation for this phenomenon. It is not merely that the philosophical novice lacks the sensitivity to properly account for the morally salient features of the practice, it is also that what moral contribution those features make cannot be expressed in abstract, tractable, propositional form. Whether and to what extent a feature counts in favor of taking some action or holding some propositional attitude is, on Dancy's version of the anti-codifiability thesis, always dependent on a holistic evaluation of those grounding features in context—never available in

general, tractable propositional form. This means that even when the philosophical novice *does* develop to the point where their grasp of the discipline improves dramatically, more information may come to light that leaves the previous information undisturbed and yet, when combined with it, drastically changes the nature or significance of a candidate consideration. In short, Dancy sees it as no surprise that the philosophical novice does not find the reasons to study philosophy ready-to-hand, because she does not yet possess the relevant sensitivity to see them. Moreover, such reasons do not hold still; they are functions of a robust sensitivity that she lacks and are not expressible in abstract, tractable, propositional form outside of their holistic evaluation. As one's sensitivity develops, the reasons the agent is in a position to see shift and sharpen in turn.

As a result, Dancy's view seems well prepared to explain why the philosophical novice might then be unable to justify her study to others. Her reasons for studying philosophy are, indeed, at this early stage, unavailable for ready articulation or appreciation by the agent herself. This is because she lacks the necessary sensitivity to attend to all the features of philosophical study that mark its practice as morally relevant. What's more, even if she did possess this sensitivity, considerations are reason-giving only as a function of their holistic evaluation, which transcends codification in tractable, abstract, propositional form. So, even if the novice could appreciate them, they resist being *given*, once and for all, to all comers. Full justification in this case is not possible because the requisite experience is lacking and the expectations for what counts as justification (the codification of rules for when something is a genuine, relevant reason) cannot be met.

Dancy's view seems well-positioned to explain our example of loving relationships as well. We saw that the central issue with the standard methods of representing the good in this case was that they each required an articulacy that is simply unavailable to those outside the

relationship, or indeed those within. This is because the value of any one particular person lies in the progressive unfolding of a historically situated life together, one whose value cannot be fairly stated in advance of first-hand practical experience with that person. Particularists like Dancy are happy to grant this on two accounts. Since holism includes a demanding context sensitivity, Dancy can agree that no one, including the agent embarking on a loving relationship, is yet in a position to say why one loves *this* person. This is, in part, because the work of loving relationships is precisely building one's familiarity and understanding of the complex details of the person who is of central concern. Without the relevant direct practical experience, Dancy can side with us in holding that the values of love cannot be fairly represented in general, absent the relevant experience needed to see them.

Another route to this same conclusion requires Dancy to lean heavily on the thesis of anti-codifiability. On this view, Dancy might agree once again with our conclusion that the value of particular human beings is not available for representation as reasons in the realm of interpersonal justification. He might explain why by pointing to the fact that features of individual human beings, like other features of the world, do not always ground the same moral meaning. What grounds my love for my partner's sense of humor is not identical to someone else's love of *their* partner's humor, even if the general explanation remains the same. Each individual person, while both funny, have their own unique style, delivery and wit. On Dancy's view, the moral contribution that the character trait *funny* makes on each occasion is itself dependent, in a way that escapes useful or tractable articulation, on what other nonmoral features are present or absent, considered holistically together. As a result, anti-codifiability leads us again back to the same place: the value of individual human beings resist being brought into the form

of publicly shareable reasons precisely because their status as genuine reasons cannot be reified into a form suitable for all to see.

With these examples in view, Dancy's particularism looks poised to solve our problem by providing its own unique version of the substantive view that we considered in the previous chapter. Dancy could take the substantive approach by emphasizing a commitment to radical context sensitivity. Which goods count as morally relevant is always dependent upon a rich contextual understanding of all the features of the case. While this context changes from case to case, it is still the case that the ideal moral agent is one who successfully sifts through those rich, complex details to identify that which is morally relevant here and now, and that which isn't. Again, the particularist who takes this approach could argue that making substantive claims about which considerations are and are not morally relevant, though they will not hold generally across cases, is constitutive of moral deliberation.[76]

IV

Dancy's particularism appears to offer some enticing answers to our worry about reasons. Specifically, holism about reasons seems to avoid the implausible generalization across novel situations that seemed to create a problem in chapter one. Since the status of a consideration depends upon a deep understanding of the situation, a Dancy-style particularist will not be tempted to explain why we pursue philosophy, nature, art or human relationships in hollow, general terms. Moreover, since Dancy holds that this context dependence is not codifiable, at least not in terms that are tractable and useful for moral deliberation and justification, it seems as

[76] It might seem as though a particularist is less likely to sign up for a version of the formal approach because the anti-codifiability thesis is essentially a denial of the possibility of establishing general formal conditions for what counts as a reason. However, as we shall see in the following section, looks can be deceiving.

though he is also happy to leave open our explanation of what reasons we have in cases where complete specification of the particular situation is beyond the pale. All this marks Dancy's particularism as a potentially appealing alternative.

But looks can be deceiving. Let us look closer at how Dancy conceives of practical thinking. On his view, "[w]hen someone does something, there will (normally) be some considerations in the light of which he acted—the reasons for which he did what he did. There are not so many things that we do for no reason at all." Indeed, he thinks this so obvious a point that he opens his book with it, calling the claim "a perfectly ordinary and unproblematic distinction, expressed in terms of reasons."[77] Our discussion in chapter one began with a set of counterexamples to Dancy's perfectly straightforward distinction—cases where it looks as though, by his lights, we *do* act for no reason at all; but not in the absentminded way one scratches an itch. Instead, these were activities whose value lies outside the linguistic purview of intelligible reasons. Dancy overlooks these cases. Why?

One explanation is that his vision of practical thought in moral deliberation is fairly narrow: "[i]f I have a duty to do this [or any] action, there must be some feature of the situation that makes it so. If in some general way I ought to do it, still there must be some feature of the situation that makes it so."[78] There are two aspects of this view worth mentioning. The first is that Dancy assumes that reason-giving features of the world can always be meaningfully individuated from the surrounding context, which gives them their normative force. Call this the *individuation* thesis.[79] The second aspect is that, once recognized as such, these individual features of a given situation generate reasons that can unproblematically serve as the basis of interpersonal

[77] Dancy (2000), 1.
[78] Dancy (2004), 33.
[79] This term is from Tsu (2018).

justification. After all, if general principles do not generate duties (because they do not exist), and we agree that duties can and do exist, then what else is left to generate them?

All this might sound straightforward, but Dancy is not clear about the nature and role of *features*. What, on his view, is being said when a feature of the world is isolated and said to provide *the* reason for which one acts? After all, Dancy is committed to the claim that a feature's status as a reason is not straightforward in the sense that it is guaranteed by any general principle that tells us what reasons are given by what circumstances. It's only through the holistic deployment of a developed practical judgment that one can see whether a reason exists here in this particular array of features that make up a given situation. Seeing in this sense is not capturable by any general principle. The question then becomes: on Dancy's view, is it a feature *itself* that generates duties or is that label shorthand for a vast set of circumstances that collectively explain something?

Suppose I return some garden shears that I lent from a friend, citing the fact that I promised my friend that I would return it by today. Is the fact that I promised *alone* what generates my reason? According to holism, whether my promise provides me with a reason to return the shears is dependent upon a holistic evaluation of the circumstances: have I finished my pruning? How important is it to my friend that they get their shears back on time? Have I encountered some family emergency that keeps me away from home? According to holism, the status of my promise as a reason is not grounded in some single feature or other, taken in isolation, but rather a holistic evaluation of the entire situation.

The point is that a whole circumstance must be in place in order for a single feature to fix its status as a reason; it is only when all these features are in place that this feature "promising" can be recognized as a reason. Take away any one part of the situation, according to holism, and

things could change. It seems that we have lit upon a tension within Dancy's view. The tension seems to sit between holism and the ease with which everyday practical thought seems to identify individual features that generate reasons.

Dancy might respond that when I speak of my promise as *the* reason for returning the shears, I am citing a heuristic that helps focus an evaluation of the whole circumstance. When we speak of "promising" as the reason for returning a loan, we are speaking in heuristic code—when we say that *this* feature generates a reason for action, we really mean something else. Importantly, he might say, nothing about this shorthand suggests any principles for identifying reasons. Nothing about promising (in general) suggests that, whenever we have a promise, we have *always* end up with a reason to act. Not every promise must be kept, come what may. When I cite a promise, I am picking out a feature of a circumstance that calls for focused attention, but it remains true that promises only count as good justification if a lot of other things are already in place: the promise was not coerced, the ask was reasonable, no harm was done, etc. The point is that many conditional requirements are assumed to be in place when we speak of a single feature as a reason. This or that feature of a situation can be a heuristic shortcut to focus in on a reason, but on one plausible version of Dancy's view, the reason proper is always the sum total of the circumstance taken together.

Perhaps this eases the tension in a satisfying way—perhaps it doesn't. What is important for our purposes, though, is that talk of given *features* of the world, combined with an insistence that they play a central role in generating individuated reasons for action, sounds quite familiar—specifically, it prompts comparison to the role that reasons play in a standard generalist account of practical thinking. So let us tug on this thread.

For starters, both Dancy and the generalist roughly agree about the broad strokes of practical reasoning. Dancy holds that "our reasons are *given* to us by features of the situation, rather than by our own psychological states—unless those states function merely as features of the situation. Our reasons, as I have put it on occasions, are objective rather than subjective or relative to our psychology."[80] While inextricably tied to features of the particular situation of the agent, Dancy is committed to the idea that reasons are independent from the requisite actions in the sense that the former can be encountered and understood independent of the latter.

The generalist sees things similarly. For example, while decidedly opposed to particularism, T.M. Scanlon's account of reasons signs up for the conceptual separability of reasons and actions. For him, reasons are considerations that reflect features of the world, which count in favor of taking judgment sensitive attitudes toward propositions or actions. Like Dancy, Scanlon thinks that these considerations are not dependent for their existence on some feature of the agent (again, unless these states function merely as features of the situation), but are rather reflective of the reality of the situation that confronts them. "If I take myself to have reason to do something because it is worthwhile—to work to alleviate some people's suffering, for example, or to prevent the destruction of some great building—this reason does not seem to depend on my seeing it as a reason."[81] Just as "the thinking of a good mathematical reasoner 'represents' or 'tracks' the truth about arithmetic insofar as it takes into account the right considerations in the right way [...] Similarly, in order for judgments about reasons to be taken to be about some subject matter independent of us in the sense required for it to be possible for us to be mistaken about them, what is necessary is for there to be standards for arriving at conclusions about

[80] Dancy (2000), 69 italics mine.
[81] Scanlon (1998), 42.

reasons."[82] While they disagree about the role of reasons in justification, Dancy and Scanlon (particularist and generalist) agree that practical thought exhaustively consists in actions chosen on the strength of reasons, which are understood as facts about the situation.

Pushing further, Dancy is even open to the possibility that there are *invariant* reasons that do not vary as a function of their context at all. Dancy concludes, for instance, that "particularism should accept the possibility of invariant reasons, so long as the invariance is not a matter of the logic of such reasons, but more the rather peculiar fact that some reasons happen to contribute in ways that are not affected by other features."[83] Interestingly, invariability is one of the central features of an atomist account of reasons. For both the particularist and the generalist, then, justification works through the elucidation of relevant features of the world. What's more, on both accounts, these features exist and are apprehended prior to (and subsequently count in favor of) an action. Moreover, some of these features can remain unaffected by changes in other features in the world.

Dancy and generalists also broadly agree about the relation between reasons, actions and justification. For Dancy, to ask for a justifying reason "asks for an account of whether the action is worth doing."[84] Once again, this looks very similar to Scanlon's account. On his view, when we encounter features of the world that could serve as reasons, "one tries to see whether this reason would be included in the most coherent and complete account of what reasons there are."[85] This process of determining what does and does not count as a genuine reason is all there is to his

---

[82] *Ibid*., 63.
[83] Dancy (2000), 136-137. It is difficult to see what is meant by the ""logic of such reasons." Dancy seems to have in mind here reasons that are somehow isolated from the existence of others. Perhaps he has in mind something like absolute prohibitions on, say, murder or rape. After all, what makes our reasons to avoid such acts so strong is precisely that they do not seem affected by changes in circumstance or the presence of other reasons.
[84] Dancy (2000), 8.
[85] Scanlon (1998), 68; Scanlon (2014), 115.

picture of practical thought—it begins with a picture of the world as a set of individuated features that call out for consideration under a certain description. For both the generalist and the particularist, normative conclusions about moral right and wrong depend exclusively on conclusions about what individuals in certain circumstances have reason to want and to do, which are to be given by features of the situation. The task in interpersonal justification is to articulate considerations whose content sheds some light on what is worth pursuing in the action as part of the most coherent and complete picture of the world.

If this is right so far, it opens up toward yet another shared assumption between the Dancy-style particularist and the generalist: that the point and final result of practical thinking is the issuance of an action, taken to be the locus of moral life and a critical juncture in the assignment of moral responsibility. Both assume that actions are performances produced by an agent on the strength of the agent having affirmed the basis of their supporting reasons.[86] While Dancy denies that the conclusions of deliberation can be universalized and applied to other cases, both agree that in order to qualify as an action at all, an act must be intentionally chosen for a reason. To do so is to place the act under a certain description that highlights what its justifying reason counts in favor of. While they disagree about whether every episode of practical thinking commits one to a general principle or how reasons are to be assessed across disparate circumstances, they share a preconception of the shape that behavior must take in order to properly count as action, that is, the relation between reasons and actions. Taken together, all these assumptions for a shared conception of the *shape of practical thought*.

[86] Korsgaard holds the clearest version of this view. For her, an act is a bit of behavior, like returning my garden shears. An action includes the reason for which it is done: returning the shears on the strength of my promise. In short, an action is a whole performance whose full description includes the practical thinking that shapes it.

What is that shape? Dancy tells us that the task of practical thinking is "an attempt to capture in the mind a shape that the situation has independently of whether we recognize it or not. And that shape is practical, since it consists in a configuration of considerations relevant to one's choice of action."[87] He continues, "The shape of our thought is simply the way we shape the situation up, but the situation has a shape of its own, which we are trying to get our thought to fit."[88] This implies that the work of creating the relevant values to be recognized in the world is done before the agent arrives on the scene. Indeed, this conception of practical thinking holds that "what we are trying to do is to establish what reasons are present in the case before us."[89] According to Dancy, "when I reason to action, my reasoning serves up an action-description, and I act accordingly." The result is that, "[w]hen an agent deliberates well and then acts accordingly, the action done is of the sort most favoured by the considerations rehearsed, taken as a whole—just as when an agent reasons well and then believes accordingly, the belief formed [or action performed] is of the sort most favoured by the considerations rehearsed."[90] Here, then, is Dancy's gloss on practical thinking: the agent first observes the world as a landscape of considerations that each count in favor of certain actions; these considerations exist independently and are conceptually prior to the agent's perception. Next, the agent considers which action is most favored, weighs and considers the result and acts accordingly.

[87] Dancy (2018), 3.
[88] *Ibid*., 3.
[89] Dancy (2004), 142.
[90] Dancy (2018),29.

Dancy's account looks strikingly similar to the generalist view.[91] For Scanlon, this picture of practical thinking is most evident in cases where there is some question or conflict about the status of a reason: is it *really* a reason for action? Scanlon outlines four stages.

> The first is the one at which X seems to be a reason for attitude A (or seems not to be). Second, there is what might be called the first critical stage, at which I may decide, as I have just said, whether X really is a reason for A or not—for example, whether the pleasure of eating ice cream is a *pro tanto* reason for intending to get some, or whether it is no reason at all. [...] Third, there is the second critical stage, at which I decide whether, taking account of X and of whatever other reasons I take to bear on the matter, there is sufficient reason for adopting A. Finally, fourth, I may come to have attitude A—for example, I form an intention to go in search of ice cream or to buy a computer. (Scanlon 2014, 102)

Scanlon's account of how we apprehend reasons grounds his view of what reasons are. He outlines a four-place relation: a reason is a relation R(p, x, c, a) holding between a fact, p, an agent x, a set of conditions c, and an action or attitude a. Compare this to Dancy's view. Officially, Dancy's account of favoring has only three relations, but he is quick to wave this aside as a minor difference of expression.

[91] However, he is careful to note that "the domain of practical reasons is not a unified subject matter like the domain of sets, the content of which we should expect to be determined by overall principles characterizing this domain. [[...] our confidence (often justified) is a matter of confidence in those particular conclusions rather than a general confidence that all questions about reasons have determinate answers, whether we have reached them or not." (2014, 104)

> I am just going to assert that Scanlon's extra place is not necessary: the conditions for the obtaining of a relation will not all of them be part of that relation. So we can allow that what favours what may vary according to the nature of the situation at issue, without allowing that this means that the general situation occupies a place in that relation similar to that occupied by the agent, the consideration and the type of response favoured. (Dancy 2018, 36.)

Dancy is in broad agreement with Scanlon here. For Dancy, the situation in which a reason occurs plays a significant role in determining which action(s) the reason favors, but we need not assign that situation a formal role in the favoring relation. Indeed, he goes so far in proclaiming his adherence with the generalist that he says, "People do have different views about the favouring relation, but they pretty much all allow that some such relation exists, and that it has (at least) the three places I have listed."[92] The point, for our purposes, is that even Dancy is willing to admit that he is in broad agreement with the generalist about the formal structure of how reasons count in favor in practical thought.

The problem here is that, as we have seen in the prior chapter, there are important cases where the consideration in question cannot be conceptually separated from the attempt to engage wholly in the activity through which it is revealed, and so talk of reasons may not seem so central to all practical thinking as Dancy and Scanlon suppose.[93]

---

[92] Dancy (2018), 36.

[93] Scanlon comes close to acknowledging this when he claims (2017) that the existence conditions of normative facts cannot be settled independent of their use in deliberative context. However, this is a separate point. Scanlon's point is intended as a defense of a robust realism about normativity that comes with its own set of ontological questions that have thus far gone answered. My point is not an ontological one but an epistemic one: what reasons there are (however they come into existence and whatever their nature) cannot be appreciated from outside the practice they are meant to inform.

Nor is it the case that just because one fails to articulate the value of some activity, experience or relationship—perhaps because one is not yet in a position to fully appreciate its value—that there is no motivation to pursue it to be found. For, again as we have seen, this would mean that there is no motivation at all to undertake pursuits whose value is not able to be seen from outside their practice. Since the unique values that those pursuits reveal cannot be fully appreciated absent first personal engagement with the practice in which they occur, appreciation of the relevant considerations is necessarily tied to some practical experience.

Whatever their ontological status—which as we briefly saw in chapter one constitutes a wide and contested literature unto itself—why we pursue these centrally important human activities is hidden from view until long after first contact with the activity, experience or relationship.[94] Such goods, whether they are generated or merely discovered by full throated participation in the relevant activity, are not conceptually distinct from the actions required for their appreciation. Nor is their value necessarily visible, ready to be given or otherwise available for articulation, even by the agent who finds themselves in their gravity. Instead, they are only revealed stepwise through progressive attempts to bring them into view. Dancy's picture of practical thought, like the generalist, leaves little room for this way of valuing. This is because Dancy places the favoring relation at the center of his account. Indeed, practical reasoning, on his view, just is the process that takes us from the situation as we see it to acting in the way that is most favored by that situation.[95]

Despite initial appearances, then, Dancy's particularism shares quite a few foundational assumptions about reasons and practical thinking with the generalist. Both agree that reasons are

[94] Peter Shiu-Hwa Tsu (2018) calls reasons that can't be individuated from their contextual enabling conditions "embedded".
[95] Dancy (2018), 42.

conceptually separable from, and prior to the attitudes and actions they are meant to count in favor of. They also agree broadly about the way reasons interact with practical thought in a two-stage affair. First, the agent apprehends what are taken to be neutrally describable features of the world, and next they connect these up to some evaluative criteria that elevates those features to the status of a reason (be it a general principle or a holistic evaluation). Finally, since both take the consideration of reasons to be the essence of practical thought, neither can make sense of the mode of valuing highlighted in chapter one.

As a result, the explanations available to the Dancy-style particularist for how we distinguish between genuine and apparent reasons do not substantively differ from the generalist either. Since Dancy holds on to the assumptions that (i) reasons are the central operative unit of practical thought, (ii) that they are independent from the actions and attitudes they purport to count in favor of, and (iii) that reasons are objectively given by features of the world, the backdrop of his picture of reasons in practical thought is shared with his generalist opponents. Both affirm the structure of practical thinking that involves apprehending justificatory features of circumstances and action types. My claim is that practical thought includes unfolding values that are internal to an ongoing activity and cannot be offered as reasons for action.

To clarify, consider one more example. Suppose I am drawn into the gravity of love, but lack the words to articulate the source of my affection. By emphasizing their commitment to a demanding context sensitivity, the particularist could say that it is quite unsurprising that I lack words to articulate the full force and form of the value I see before me. On the particularist view, which considerations count as practically relevant is always dependent upon a rich contextual understanding of all the features of the case. Perhaps I lack an informed sensitivity to the well of history and idiosyncratic details of my beloved, in which case I would subsequently lack the

words to express those details that I miss. While this context changes from case to case, it is still the case that the ideal practical thinker is one who successfully sifts through those rich, complex details to identify that which is practically relevant here and now, and that which is not. Again, the particularist who takes this approach essentially makes substantive claims about which considerations are and are not practically relevant, though they will not hold generally across cases, and holds that this is constitutive of practical thought.

Dancy's view has appealing features. Let me attempt to reconstruct what he might say about love. For one, the phenomenology of embarking on a loving relationship includes a deepening understanding of each party as they move through time, sharing more and more of life together. But this picture has two fatal features. The first is its commitment to the goods that make actions worth going for as *given*. In practical reasoning, "[w]e are trying to get it right, and getting it right means conceiving of the situation as in fact it is, in all its glorious complexity. So there is a shape to the situation and we are trying to get our thought to fit that shape. The shape of the thinking is intended to match the shape of the situation."[96] His brand of realism takes reality to have what he calls a "normative shape" which comprises the ways in which the various considerations that exist in it combine to call for one form of response. Dancy assumes that those genuine reasons exist, ready-made to be seen from the start, as identifiable, static, or indifferent entities to the changing individuals who make up that relationship. Where Dancy's view finds the value of an activity in the circumstances, whose features provide a reason to act here and now, I am highlighting the value of interactivity itself.

Surely, what is worth going for in a loving relationship can include emergent properties that from sharing a life together. This emergent good cannot always be stated or seen from the

[96] Dancy (2018), 3.

outside because it does not exist absent the direct attempt to bring it into view. The value of this good, as we have seen, is not available in the form required by philosophical theory to be offered as justification. My point has been that it is this theoretical philosophical bias, and not the lived experience, the phenomenology of love, that stands in our way of understanding what goods are lit by the light of love.

Dancy's notion of practical thinking is essentially the same as the generalist. Both agree that what is morally relevant and worthwhile can be contained within the space of reasons. Both also roughly agree about how reasons count in favor. Thus, Dancy's account of reasons in practical thought does not so much answer our worry (about the unjust exclusion of a range of human goods from the domain of practical thought) as push the problem one step back. While Dancy has a unique response to the methodological problems that flow from an insistence on unchanging general principles as the source of normative authority, his view fails to respond to our methodological concern about the role of the corresponding assumptions about reasons in ethics. This is because he, like the generalist, takes on the very same grounding assumptions about what reasons do in practical thought and how we encounter them in the world. Since Dancy shares these key foundational assumptions about how considerations must be given and formulated in order to serve their dual purpose, his approach does not have the resources to respond to our exclusion concern. Absent a radical revision of these methodological assumptions, an answer to our present concerns does not seem forthcoming.

V

While he differs dramatically in his respective account of how and when reasons count in favor, Dancy's particularism takes on a fairly standard vision of the relation between reasons,

action and practical thought. Both he and the generalist see reasons as an independent feature of the world, as *given* to us, as fundamentally separable from the actions and attitudes they purport to justify. Both assume a very high standard of articulacy in their understanding of how reasons function in interpersonal justification. Second, Dancy shares essentially the same view of what reasons do (how they function) in theories of individual moral deliberation and interpersonal justification, and their centrality in practical thought. Dancy's contentful revision of moral theory does not provide any new resources to answer our concerns from chapter one.[97] In short, it shows us that the root of the issue of unjustly excluding an underexplored class of considerations that properly features as reasons in moral deliberation, the problem identified in chapter one, is not the content of any particular generalist theory, nor the generalizing impulse itself.

Instead, the heart of the issue is a set of widespread methodological commitments about practical thinking. The first is that reasons are considerations that flow from brute features of the world, neutrally describable to all comers. The second is that such a view of reasons is the entirety of practical thought. The issues we raised in chapter one about reasons connect to action, how we encounter them, and the conception of agency that takes reasons as its central currency, combine to artificially limit the scope of reasons in models of practical thought by the particularist as well.

This limitation also represents a missed opportunity. After all, we have seen how holism about reasons and anti-codifiability seem to produce a clearer picture of how deliberation works

[97] At this stage, it could be suggested that what is really at issue here is the legitimacy of the moral realist project. For instance, one might be tempted to interpret my claim from chapter one about appreciating the value of philosophical study, art or individual relationships as really being that in these cases talk of reasons is out of place because *there are no reasons* absent the first-personal attempt to create them. For a good summary of this debate, see McNaughton (1988), especially chapter 6. While this may be a related question, the reader will recall from the introduction that my argument is neutral about the status of reasons. Instead, my focus is on how they are deployed in moral philosophy, their functions and methodological use, be they facts in the world or constructions of the mind.

in the cases we identified. Where the clear articulation of reasons does play a central role in most instances of moral deliberation, we saw in chapter one that reliance on reasons talk cannot account for cases where appreciating the value of a person, experience or activity is inseparable from the first-personal attempt to bring that value into view. Holism about reasons offers an explanation of why articulation of what matters in these cases is at best a meager approximation: because full understanding of the moral character of the consideration in question can only be had through a holistic evaluation of the context in which it occurs. Moreover, allowing that the ways in which a consideration functions within that holistic evaluation cannot be codified in tractable rules explains why the articulation of the value one initially saw in such activities rings hollow. This leads us to a surprising conclusion: the phenomena we highlighted from chapter one—studying philosophy, appreciating art and loving individual people—seem easier to explain on the particularist view. As such, their exploration seems ripe to serve as compelling objections to generalism. Yet what we have seen in this chapter is that Dancy (particularism's most prominent advocate) not only misses this opportunity, but cannot himself account for these centrally important phenomena.

We need a new path forward. This means creating a particularism of sorts that can be severed from the philosophical prejudices about practical thinking that plague Dancy-style views. To preview how we might chart that path, consider one final example. Suppose I am attempting to play the cello. Excellence in that activity, attempting to play with sufficiently rich tone, in tune and tempo, with correct dynamics, is what I have a glimmer of as I begin. A glimmer of that excellence is needed in order to be moved to engage in its pursuit at all. That glimmer guides my activity, but not in terms of reasons to do it one way or another. In cases like this, talk of reasons

does not fairly describe this practical guidance, the kind that can only be artificially and with distortion translated into the language of reasons. We also have a direct apprehension of goodness: excellence is not a relation between one thing and a thing that counts in favor, it's seeing one thing: how to play. Contemporary philosophical theory that we have so far discussed takes what is in fact *part* of practical thought, careful consideration of reasons, as its whole.[98] Talk of actions intentionally done for reasons fails to light up what is valued in that incremental, ongoing form of practical thought in musical study. Developing a view that can is the task for the following chapter.

[98] Indeed, Scanlon says nearly that: "it seems to me that such a process of careful reflection is the only way we have of arriving at conclusions about reasons for action." (2014, 102)

Chapter 3: Rethinking Practical Thinking

I

Our investigation began by identifying a deficiency in the way philosophers, by and large, theorize about reasons. We saw that the way individuals value some important activities, experiences and relationships resist formulation in the language of reasons. We then traced this deficiency to a set of grounding assumptions about the structure of practical thought in philosophical theories. We uncovered a limited picture of practical thinking that relies almost entirely on talk of individuated reasons, which included an overly rigid account of how we encounter such considerations and mark them as sufficient for justification.

Chapter two attempted to sidestep these assumptions by turning to Dancy's particularism. Though initially promising, upon closer investigation, we saw that particularism's most prominent advocate actually shares these grounding assumptions about practical thought that we identified as the root of the problem. What is needed, we saw, is a fresh way of looking at the relation between valuing and practical thought.

This chapter explores the philosophical work of Iris Murdoch, asking what resources her account of practical thinking might offer. In particular, it interprets Murdoch's early work on love, attention, imagination and the experience of concept acquisition to illustrate how philosophical inquiry can better capture the modes of valuing we discussed in chapter one. By considering and then sharpening these terms, this chapter lays the groundwork for my positive account of practical thought in the final chapter.

As we will see, the breadth of Murdoch’s philosophical work brings with it some intimidating and unresolved questions. But if I can clarify the handful of concepts that seem

central to understanding her view of practical thinking, and in so doing set the stage to apply these concepts to my own account in chapter four, the present discussion will be a success.

II

Iris Murdoch's philosophical work has recently enjoyed a well-deserved resurgence.[99] One prominent locus of discussion is whether or not that work supports a version of particularism. Some claim that Murdoch was a pioneer of the approach.[100] Others hold that Murdoch's work supports only a certain brand of particularism; perhaps she was a particularist in method but not substance.[101] Still others contend that Murdoch was not a particularist of any sort.[102] This disagreement centers a growing literature whose unifying goal is to reveal, once and for all, whether Murdoch was a particularist.

I think that the focus of this literature largely misses her most momentous contribution to the field—a reconceptualization of practical thinking.[103] Like the particularist, Murdoch critiques the usefulness of general principles and champions the importance of attention to particulars. But unlike the particularist, her philosophical work largely rejects the standard view of practical thought that grounds most mainstream philosophical theories. Murdoch rejects what she sees as modern philosophy's obsession with rationality, the scientific flavor of modern theorizing, its corollary conception of the world as given, and the tight relation it asserts between reasons and actions.

---

[99] See Panizza (2022; 2019), Denham (2012; 2001), Crisp (2012; 2000) Millgram (2005), (Little (2000).
[100] Blum (1994), especially chapters 1 and 3.
[101] Driver, (2012), 293-305.
[102] Hopwood (2018) focuses primarily on Murdoch's fiction to widen the scope of her philosophical import.
[103] The exception is Diamond (1996). I will discuss her view in the following section, and more extensively in chapter 4.

Consider Barbara Herman's picture of the role of reasons in practical thought. For Herman, a consideration's status as a reason for an agent depends on the development of a capacity she calls *moral literacy*. Moral literacy is "a capacity to read and respond to the basic elements of the moral world."[104] This capacity begins with an acknowledgement of certain basic facts such as the difference between persons and objects, the special status of the latter in virtue of their capacity for autonomous action and the culturally fixed lexicon of wrongs and injuries that any morally sensitive person can be expected to recognize. The development of the capacity for moral literacy is the development of a practical disposition that enables the reliable recognition of morally salient basic features of a circumstance and action, as well as the regular capacity to do what is seen to be right."[105] On this view, the recognition of a consideration as a morally relevant reason is dependent on a capacity for *reading* the morally salient features of the situation. Competent development of this practical disposition is a form of literacy that marks a normative standard for agents; it is a standard for mastery over an ability to consistently grasp the moral meaning of features of the world, much in the way reading literacy marks a mastery over one's ability to grasp the relevant meaning and salience of a written language.

While the content of her Kant-inspired account of morality that follows is disputed, Herman's view of practical thinking is instructive because it depicts a rather familiar center of gravity in the literature. For views like Herman's, practical thinking essentially consists in (1) responding to reasons, which are (2) grounded in neutrally describable facts in the world, which

[104] Herman (2007), 97. Herman has since said much more about the elements of that "moral world." For Herman (2021, 181) it is a *moral habitat*, comprising a "descending matrix of duties, juridical and ethical, perfect and imperfect, extended by entanglement, adjusted by casuistry, and supplemented by the duty of due care, is a complex and dynamic system."
[105] *Ibid*., 98.

(3) exist in some sense independent of our perception of them and (4) bear relation to our thoughts and actions by counting in favor.

Murdoch sketches a radically different picture. Her account is *bifurcated*. When it goes well, practical thinking is a gradual, infinitely perfectible task, which involves both overcoming one's self—our initially limited set of concepts—to develop a heightened sensibility to what is valuable in the world. When it goes poorly, (when we fail to see the world as it really is) practical thought involves a failure to *attend* properly to the object of value, which results in a flight into *fantasy*, where the agent becomes insensitive to the different dimensions of normative salience that are appropriate to the object of thought. On her view, moral *concept acquisition*, *imagination*, *attention*, and *love* are all rich, important aspects of how agents move from thought into practice. With even these preliminary remarks, we can see that Murdoch's account of practical thinking opens onto a much wider conceptual field than the accounts we have so far discussed. At the very least, thinking with Murdoch will take us far beyond reasons for action.

Figure 1: Murdoch's view of Practical Thinking

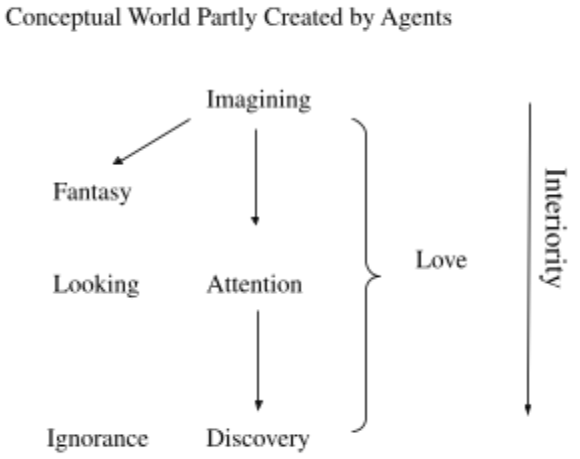

These are provocative claims. But our initial gloss conceals an uneasy tension. By the time we begin to deliberate at all, Murdoch thinks we choose between restricted options in a world that we have already partly created.[106] Which features each of us are in a position to see, and to what extent, which options are alive for us and which unthinkable, will be partly a function of the normative concepts we are able to grasp at the time. At times, Murdoch speaks seriously of agents "living in different worlds."[107] Yet *other* times she speaks of practical thinking as an act of *discovery*, as a peeling back of the veil of our own conceptual blind spots to see the world as it really is. So, which is it? Does Murdoch think practical thinking is an act of creation or discovery? Murdoch is less than clear. A careful reconstruction is in order.

Murdoch's view of practical thinking begins with an account of what moral concepts are and how they shape practical thought. On her view, cognition is always a moral activity.[108] Thinking, and so the acquisition of concepts, are activities we do as inescapably moral beings. The root of this claim is a rejection of what has become known as the fact/value divide. In the 1950s, philosophers in the analytic tradition were, by and large, deeply committed to the position that no evaluative conclusion follows from purely factual premises. The thought behind dividing facts and values is this: when presented with the facts of a situation, a moral agent can always describe them in non-moral words. He then can then express his commitment to an evaluative principle or framework by exercising his freedom to choose criteria for the use of such terms as *right* or *good*. The appeal of this fact/value divide was allegedly to protect against the dangers of the naturalistic fallacy: arguments that move from factual premises to evaluative conclusions

[106] Murdoch (1999), 199-200.
[107] 1999d), 81-82.
[108] Murdoch (1992), 26.

depend, so the thought goes, on some (often suppressed) attempt to define moral terms in nonmoral terms.[109]

But one might rightly ask, *what is the real danger here*? Cora Diamond provides a window into Murdoch's response to this type of approach to dividing up the world: "what is the matter with suppressing a premise in a moral argument may be that the hearer of the argument will not grasp the possibilities open to him, will think that he is forced to the conclusion by undeniable facts."[110] This might be a problem, for instance, if the thinker is attempting to do something they cannot square with their conscience. Thinking about why philosophers so strongly desire to expose a suppressed moral premise in the first place, Diamond points out, connects voracious support for the fact/value divide to a set of deeper moral commitments: in particular, the desire to take responsibility for our actions and beliefs. Where the naturalistic fallacy suggests that some undeniable facts of a case can fix the moral rightness of its conclusion, dividing facts and values precludes this possibility: only *human beings* can fix the moral rightness of a conclusion by freely choosing some evaluative criteria or other. As a function of this choice, we are responsible for affirming the conclusion of an evaluative argument.

Murdoch lights upon this underlying commitment to responsibility in moral decision making, but reaches the opposite conclusion: we cannot cleanly separate facts and values. For her, it *is* important to expose suppressed premises in arguments, but not for any purely logical reason, as supporters of the fact/value divide claim. Instead, she points out that the motivation for holding fast to the fact/value divide as a matter of logical tidiness reveals a deeper moral sensitivity, a recognition of the possibility of the evasion of personal responsibility through

[109] *Ibid*., 53.
[110] Diamond (1996), 81.

argument. The real problem with the naturalistic fallacy is that the agent thinks the *world* has decided the moral rightness of their act, not *them*. The agent might claim to be merely "following the facts." We ought to resist this claim. But *why* we ought to resist it, Murdoch suggests, also gives us reason to deny the fact/value divide. The line between facts and values, then, is not merely a matter of logical clarity, but is also indicative of a deeper moral stance: a hidden derision of attempts to evade moral responsibility for our choice of evaluative criteria.

Murdoch's insight is this: denying the legitimacy of deriving moral claims from factual ones is *itself* a moral plea: to take responsibility for our evaluative attitudes, for *how* we choose to view the world. This suggests, for Murdoch, that disagreements about what falls under the categories of *fact* and *value* are unlikely to be settled by logical analysis. Instead, Murdoch puts forward a different characterization of practical thought, one that rejects the fact/value divide, where most concepts are inherently morally neutral and only gain evaluative status when brought together under some chosen evaluative criteria.

In its place, Murdoch claims that all philosophy, indeed all *cognition* is moral in the sense that moral concepts may determine what we take a situation to be, what the facts of the case actually are. On her view, it is not always the case that if a moral concept were withdrawn we would be left with the same situation or the same facts. For Murdoch, morality is all-pervasive such that human consciousness is saturated with value. This, then, is the preliminary sense in which she thinks practical thinking is a *remaking of the world*—the concepts we deploy to understand the world actively reshape what we are sensitive to, what we are not, and will in part determine the facts of a case in the following sense. Murdoch does not deny that there are some cases where the facts of a case can be agreed upon by all. Her point is only that the application of

what may appear to be a neutral term will not be settled in all cases by readily ascertainable facts. Any philosophical position necessarily presupposes certain value judgements.[111]

To illustrate, suppose I am upset with a friend for interrupting me in conversation. In this case, there is no dispute about some of the facts: all agree that I was talking, then she, before I had finished. What *is* disputed is the appropriateness of the application of the term *interrupt* in the pejorative sense. Given my midwestern upbringing, the interruption of others appears naturally to me as "not to be done" in the same way an apple looks red. That interruption is wrong *looks* like a fact of the case. In contrast, my friend hails from a family where talking over each other is actually a sign of *interest* and engagement in the conversation. Thus, for her, interruption appears equally clear to her as "to be done". Murdoch's insight is helpful here: what has happened is that my friend and I have imbued differing moral valences to the same term (*interrupt*) because we come from different cultural backgrounds, and hence have differing grasps on what the concept means. In Murdoch's parlance: we come from different moral worlds.

Equally important, however, is that there *is no fact of the matter* about who is correct, no neutral concept which sits beneath our respective evaluations in virtue of which *interruption* is to be done or not. Appeal to linguistic, logical or any morally neutral criteria will fall short. What needs to happen in this case, it seems, is that each party needs to expand their respective understanding of the meaning of the concept *interrupt*, each needs to come to see how it could have either moral valence. If we are to grow together (rather than apart) my friend and I must

[111] Mark Hopwood (2022, 86) Notes that Murdoch does not merely assert, but spill quite a bit of ink to illustrate it. Here is how Hopwood glosses the importance of this thought for Murdoch: "Over and over again, throughout the chapters of MGM, she shows how philosophical theory is grounded in and structured by judgements of value. These judgements of value are usually not put forward as such, but are implicated in fundamental decisions about what questions to pursue and what phenomena to attend to. Often, they appear in the form of metaphors and images, or are even left entirely unspoken. Murdoch's point is not that there is anything illegitimate about grounding philosophical theory in judgements of value; in fact, she thinks that it is impossible to avoid."

modify our respective grasp on the concept *interrupt* to better understand each other's intentions. Such growth in our cognitive grasp on the possibilities for the meaning of a concept is an irretrievably moral growth because we have each learned something about ourselves, each other and what we ought to do in conversation together.

It might be objected that the best way to understand this case is actually by separating facts from values. On this view, my friend and I agree on the facts of what happened, but we simply have different suppressed evaluative premises. For me, interruption is *not to be done* because it is a sign that one is not listening to the other, and listening is a moral imperative. For my friend, interruption is *to be done* because it is a sign of active listening, which is a moral imperative. But Murdoch's point is that *who is correct* in this case will not be settled by appeal to any linguistic or logical rules. Moreover, understanding what is really happening here cannot be done in nonmoral terms because the center of the disagreement *is* moral. Instead, on Murdoch's view, if resolution is to come in this case, it will be from each party reshaping our respective grip on what is happening in the conversation—two people with incomplete grasps on a concept. Each of us must expand our set of possibilities for what *interruption* means. This process, Murdoch thinks, is both infinitely perfectible—since we can never be sure that we understand all possible meanings for a concept—and irreducibly moral—since the meanings we come to know involve what we ought to do.

This is the dual sense in which we are "perpetually moralists."[112] The first is that *all thinking is moral thinking*. Our grasp on the world, the concepts we deploy to describe it, are already infused with our values by the time we come to deliberate about them. This process of

[112] While often attributed to Murdoch, this quote is actually from Samuel Johnson's essay "Life of Milton." I came across this quote in Diamond (1996), but the reader can find it in Goldberg (1993).

unpacking (and sharpening) our intuitions is part and parcel of our development as practical thinkers. The second sense in which moral thinking is perpetual is that, for Murdoch, our grasp on moral concepts is infinitely perfectible—there is no limit that could be set to our understanding in advance of setting out. If I am to be a sensitive conversation partner, I cannot take it for granted that I have, once and for all, grasped *the* understanding of what *interruption* means for all. Instead, I must take a stance toward my own cognitive grip on the concept that it is always open to revision, new insight, that its full meaning is not a neutral fact that I can grasp once and for all.

We have taken a step toward easing our initial tension: Murdoch thinks that we perpetually remake the world we see through the ongoing revision of our moral concepts, which are basic to our perception of what is the case. This revision is ongoing, perfectionist and inseparable from our cognitive understanding of the world. While we can share with others many understandings, we do so through shared experiences. More commonly, we often see different things when we look upon the same scene; reality is fundamentally normative. This, Murdoch argues, is because all of our concepts, what they mean and our understanding of how and when we deploy them are partially informed by our understanding of what is morally salient. Moral concepts play a *shaping* role in practical thought.[113]

The picture she pushes against looks like this. A concept is thought to be like a circle drawn around an area of neutrally describable fact. What determines the area of fact overlaid by a

[113] Sabina Levibond (2021) Raises an interesting parallel worry for Murdoch's picture: "don't we lose rather than gain insight if we insist on finding, or purporting to find, 'moral' consciousness everywhere, rather than in some places in contrast to others?" The question is similar to ours: how do we understand the sense in which the world is both full of individual objects of value and also fully diffused with moral significance everywhere? Answering Levibond, I think, requires a closer look at the faculties of imagination and attention than is available in the literature. Filling this conceptual void will be a focus of the following sections.

descriptive concept circle are the relevant *meaning-rules*, which are held fixed. How evaluative overlay is placed will be a function of our decisions of principle, which means evaluative circles are movable. A secondarily evaluative concept like *interrupt* would, on this view, be tied to a particular area by meaning-rules, but people who rejected or disagreed with the moral views associated with the concept would be able to describe the relevant area in some other way, without commitment to any evaluation. Cora Diamond offers a helpful description of this view that both she and Murdoch rejects:

> Rational use of descriptive or evaluative terms is, one could say, "in virtue of" use. The meaning-rules for descriptive terms fix that in virtue of which they can be applied; so do the meaning-rules for "specialized" evaluative terms like 'courageous'; the features in virtue of which a non-specialized evaluative term is applied are determined through our own choice of moral principles. (Diamond 1996, 94)

Murdoch rejects not only the circle-picture of how moral concepts are deployed but also the very idea of moral disagreement that is dependent on it. Concepts like *good* and *right* and even *justice* or *honor* shape our view of life, but are not applied via merely descriptive criteria in virtue of certain features of a situation. Such terms can cast light on various situations; but *how* they will do so is not tied to any purely linguistic or logical convention.

Murdoch rejects the idea that moral concepts are *in virtue of* concepts. For her, a moral concept need not be thought of as like a (more or less) movable and extensible circle laid down to cover some area of fact; instead, it may be more like a lens for vision. Murdoch thinks such vision is not merely *had*, but achieved, through experience and development of our practical

judgment. While concepts like *interrupt* might not be integral to someone's life, lived experience might change the depth to which concepts like *love*, *honor* or *friendship* might reach in our understanding. That change marks a growth (or decay) in our understanding of what these concepts mean.[114] For Murdoch, the achieved understanding of a moral concept is not a matter of shifting to some new or modified evaluative criteria or moral principles. Instead, she thinks that the very idea of describing events, understanding the way the world is or could one day be, is *itself* a moral activity because it necessarily involves an exercise of moral concepts.

III

For Murdoch, all cognition is moral. Moral concepts are not *in virtue of* concepts, called into action in virtue of some deeper, neutral facts—they directly (yet often invisibly) shape and order our understanding of the world. It also means the process of acquiring concepts is ongoing, perfectionist, located in our individual history. Which concepts we grasp and to what extent is reflective of our own stage of moral development. But note that this preliminary step leaves our uneasy tension intact. It remains unclear *how* agents can (re)create the world they reason within and at the same time discover the world as it really exists. The next step toward easing that tension is to trace a distinction Murdoch makes between *fantasy* and *imagination*.

For Murdoch, imagining in this sense is an overcoming of one's particular biases, a dedicated, reflexive gaze that sometimes requires the acquisition of entirely new concepts in

[114] Here is Murdoch's (1992, 337) gloss: "Morality, as the ability or attempt to be good, rests upon deep areas of sensibility and creative imagination, upon removal from one state of mind to another, upon shift of attachments, upon love and respect for the contingent details of the world." Murdoch is careful to warn against any attempt to creates a "Murdochian moral theory", but we can begin to see what one might look like, in the following section, by tracing her distinction between creative imagination and egoistic fantasy as a criterion for distinguishing between veridical and non-veridical moral perception.

order to perceive the reality of the world. Murdoch notes that imagining is "an (inner) activity of the senses, a picturing and grasping, a stirring of desire." This is a "continuous detailed pictorial activity whereby [...] we make and remake the 'world' *within which* our desires and reflections move, and out of which our actions arise."[115] A person's *world* in this sense is a metaphorical stand-in for the set of possibilities for apprehending the real world. The world that we see, the way it appears to us, this is the world we can remake by refining our concepts. Following Plato, Murdoch takes imagination to be central to moral improvement: it is the "progressive destruction of false images." By *images*, she means something rather broad: she thinks we build up understandings of concepts, values, words, that help us grasp the reality of what is before us when we begin to think. Since apprehending an image in this broad sense of understanding is always an imperfect activity, imagination is the creative force that works as a mediator between our sense perception and the concepts we use to deploy them. When working well, imagination is a truth-making creative faculty of the mind. It works as a "sorting, separating, and connecting movement of the mind" and represents "our ability to picture what is quite other."[116]

For example, in politics, this means to "picture what it is like for people to be in certain situations (unemployed, persecuted, very poor)" and in moral theory it might mean "to relate axiomatic moral ideas (for instance about rights) to pragmatic and utilitarian considerations."[117] Imagination, then, is the name for a mechanism of the mind that connects, sorts and allows us to picture accurately what is really the case in the world. Imagination is particularly crucial when consideration of what *is not* the case for us, or is not yet the case is crucial. Just as composing a work of art takes effort, dedication and thoughtful intervention, so too Murdoch thinks

[115] Murdoch (1992), 325.
[116] *Ibid*., 323.
[117] *Ibid*., 323.

imagination helps us apprehend value in its proper context and complexity. In short, imagination is necessary for coming to grasp the value of anything in practical thought: "'becoming better' is a process involving an exercise and refinement of moral vocabulary and sensibility."[118] To imagine in this sense is to challenge and then reimagine the lenses through which they view the world is to imaginatively remake the world one sees in light of the virtue one hopes to embody.

To see Murdoch's notion of imagination at work, consider what is asked of us when philosophers begin a thought experiment: we use terms like *imagine*, *illustrate*, *suppose*, *consider*. These are terms that mark both an excursion into fictional space but also demand a conceptual attentiveness to the structure, content, truthfulness and accuracy of what follows. What is demanded in considering a philosophical thought experiment is precisely the careful, creative gaze toward truth that Murdoch wishes to bring out. The imagination at work in a philosophical thought experiment is seeking conceptual clarity, truth, and it sits on guard for details that fail to match the reality of the world; the imagination in thought experiments seeks to construct an overall image of both the fictional world of the example as well as, crucially, the parallel real phenomena in the real world, in service of which the example was created. Imagination in thought is a conceptual, infinitely perfectible striving toward grasping the reality of the world by constantly refining our concepts.

Fantasy, on the other hand, is a flight from reality. To indulge in fantasy in practical thought is to run from this task of overcoming, to hide in the palace of one's own limited, fictional construction of the world. Murdoch speaks here of vanity, vengeance, delusion of grandeur. Dreams of power, "can imprison the mind, impeding new understanding, new interests

[118] *Ibid*., 324.

and affections, possibilities of fruitful and virtuous action."[119] Fantasy is, for her, quite appropriately sometimes called *unimaginative*.[120]

Imagine a neighbor who never fails to connect a conversation to their hatred of the president. On Murdoch's view, such a person lives in perpetual fantasy, a conceptual prison of sorts. Their world is a place shaped by fear, where they punch at political opponents who don't really exist and build images of people that are caricatures of reality; this person finds conspiracy behind every corner and in so doing clearly moves not toward *truth*, but rather confirmation of their own prejudice.[121] When we, as their conversation partner, grow tired of their incessant logical leaps and fanciful suppositions, we are picking up on the banality of their conceptual landscape. The conceptual world, for them, is a very small place. This is indicative of Murdoch's notion of fantasy. It represents an unimaginative fixation, a failure to launch that locks the mind by allowing only a small circle of poorly understood concepts (e.g., liberal, conservative, freedom, etc.) to color the entirety of their vision.

Imagination and fantasy sit at opposing poles of a continuum of practical thinking. Consider what happens when we fail to understand one another. Or, to use Murdoch's terms, to fail to clearly see the individual before us. Murdoch tells us that "one may fail to see the individual because [...] we are ourselves sunk in a social whole which we allow uncritically to determine our reactions, or because we see each other exclusively as so determined."

---

[119] *Ibid*., 324.

[120] Here is Murdoch (1992, 334) on this point "'He lacks imagination' is often uttered as a reproach, but 'imaginative activity' is not always what is required. Imagination can 'go too far', 'become too personal', 'be an end in Itself'." She is careful to note, though, that "we must avoid the temptation to simplify the problem by general reference to the distinction between good creative imagination and fantasy imagination. Perhaps in some kinds of moral thinking imagination, in anything like its ordinary senses, is out of place."

[121] Murdoch (1999b, 201) relates freedom from fear with imaginative perception of reality: "Moral freedom, if it is to be defined at all, cannot, it seems to me, be defined without a reference to virtue. A mediocre man who achieves what he intends is not the ideal of a free man. To be free is something like this: to exist sanely without fear and to perceive what is real. I would be prepared to imply that one who perceives what is real will also act rightly."

Alternatively, “we may fail to see the individual because we are completely enclosed in a fantasy world of our own into which we try to draw things from outside, not grasping their reality and independence, making them into dream objects of our own.” To overcome this alluring fantasy, we must exercise love, which is always “an exercise of imagination." [122]

Through fantasy, one's gaze is turned inward; what is seen is only oneself and one's own desires. In contrast, imagination pierces through the veil of one's projections and desires. Thus, for Murdoch, the antidote to fantastical escapes will require an attentiveness that yields sensitivity to the idiosyncrasy of others without obscuring their reality. Imagination, when directed and honed by attention, is directed outward—it reveals the reality of other people in all their particularity.

With the contrast between imagination and fantasy, we have taken another step toward unraveling Murdoch's uneasy tension. We saw that for her, the acquisition of concepts that we deploy in practical thought is always shaped by morality. This constant shaping of concepts is the sense in which we arrive ready to deliberate in a world that we have ourselves partly made. But this conceptual world making can go well or not, as a function of whether it brings us closer to or farther from understanding the complex reality of the world. Imagination is the name for the capacity to construct an image of the world that is progressively closer to the truth. Fantasy is the name for the failure of imagination. But at this stage, one might reasonably ask: *how* does one move toward imagination and away from fantasy? To answer, we need to take up one more distinction that Murdoch offers in her account of practical thought—the distinction between *attention* and *looking*.

[122] Murdoch (1999a), 216.

IV

The imaginative, recognitional task of practical thought is carried out through the disciplined direction of *attention*. On Murdoch's view, we learn through actively attending to objects and their contexts. More than an occasional, discrete task to be done, the activity of attending is an infinite and inexhaustible task of attending to an individual reality.[123] *Looking* is the failure to attend. She tells us that,

> looking is the neutral version of attention, which is constantly building up increasingly coherent and systematic but nonetheless false visions of the world complete with accompanying normative vocabulary. Attention is the effort to counteract such states of illusion. […] The task of attention goes on all the time and at apparently empty and everyday moments we are 'looking', making those little peering efforts of imagination which have such important cumulative results. (Murdoch 1970, 43)

In the positive case, the agent is attending to the task of bringing the reality of the world into view. Attention is an instrument of perception; but recall that for Murdoch, what we perceive are fragments of the world as it really is, whose correct unification under sophisticated understandings of concepts is a matter of moral sensibility. Murdoch's vision of practical thinking goes on all the time; it is a piecemeal business of building up structures of value through our habitual conversations and relations with others—for better or worse.

Murdoch's picture of practical thinking shines brightest when we consider the difficult cases in life, the ones for which it may be unclear why we ought to pursue one course or another.

[123] *Ibid*., 23-30.

In such cases, the metaphor of vision steps in; Murdoch speaks of *attending*, straining to see.[124] If we try to pick apart a singular bit of straining to see, call it the "action" and separate it from the apprehension of some independently specifiable considerations that count in favor of doing it, we depart radically from the subject Murdoch wishes to highlight. For her, the work of practical thinking is to bring its object more clearly into view. This clarifying work is a progressive unfolding, a setting up of a field of view, or as Murdoch also describes it, a *magnetic field*. When in love, one can't always articulate what reasons one has to keep attending to the object at hand in advance of the activity of attending itself. Instead, such reasons only come into view in the ongoing attempt to bring one's beloved, in their complexity and depth, into focus. On her view,

> [w]e evaluate not only by intentions, decisions, choices, but also, and largely, by the constant quiet work of attention and imagination. The image here is not so much that of a body moving but rather of a sort of seeping of colour, or the setting up of a magnetic field. When moments of decision arrive we see and are attracted by the world we have already (partly) made. (Murdoch 1999, 199-200)

Murdoch names this practice *loving attention*. What is meant by *love*? "Love is the perception of individuals. Love is the extremely difficult realization that something other than oneself is real. Love, and so art and morals, is the discovery of reality." Such a discovery does not "crystallise out of the historical process of simple society living strictly by general rules" but rather is "an exercise of imagination" where the imagination opens up to the "[a]pprehension of something else, something particular, as existing outside us. [...] Love is the imaginative

[124] Murdoch (1970), 23.

recognition of that is respect for, this otherness."[125] For Murdoch, "the formulation of beliefs about other people often proceeds and must proceed imaginatively and under a direct pressure of will. We have to attend to people [...] This activity is, moreover, usually and often inevitably, an activity of evaluation."[126] This straining to see, the clarifying of our vision of the reality of the world through ongoing reflective activity, is what she calls *love*. Attention, in this sense, then, is the working vehicle of love.

When talking of love and attention, Murdoch's paradigmatic cases include appreciating the value of human beings and human relationships. But how do love, attention and imagination function in cases that involve valuing other kinds of objects? Consider a novice cellist. Upon first introduction to the instrument, the cello appears to the agent as totally *other,* a mere object—the agent lacks any sense of how to hold it, the proper shape of the hand, how to press the fingers on the strings just so, the weight of the arm on the bow, or even which sounds are desirable. The novice also lacks a cognitive grip on a whole set of musical concepts: *intonation*, *tempo*, *key* and *tone* form part of a cluster of concepts whose possible meanings vary by instrument and musical style. All this is foreign to her at the start. Insofar as she lacks any practical understanding of what it means to play the cello in these ways, the novice begins her study with only an *inkling* of the value of beautiful play.

But as she practices, the cellist learns what is important for accurate playing and how to focus her attention on just those key elements: where to relax her body and where to hold tension, how to translate written notation into physical movement at pace, coordinating her breath to the rhythm of the piece. By attending to her body in relation to the instrument, while

[125] Murdoch (1999a), 215-216.
[126] Murdoch (1999b), 199.

simultaneously interpreting the meaning of the written notation, the cellist's mind forms an imaginative connection between what exists in writing and what exists in sound. For her to progress as a musician is, in part, to strengthen her cognitive grip on the relevant concepts at play through a dedicated, focused attention to the practical elements of her play, which sometimes takes many years.

The *other,* equally significant part of this progressive development is broadening the imaginative possibilities of music itself. To move from a novice to an intermediate player is to link a practical, bodily competence (e.g., posture, weight distribution, finger pressure, speed etc.) to a basic understanding of a set of musical conceptual tools like *rhythm*, *tempo*, *dynamics* and *key*. But moving from an intermediate player to mastery requires a deeper appreciation of the those same basic conceptual tools: she learns to feel the swell of a well-done *crescendo*, what counts as soulful *vibrato*, the practical difference between vague tempos like *allegro* and *allegro non troppo*. The master cellist becomes able to ask musical questions that at first, she would not have understood. It is only because she has practical experience with the activity of cello playing that she is in a position to even be puzzled by an *interpretation* of Bach: such a phrase means little to the novice, even less to the nonmusician. Part of what marks practical mastery over a musical instrument is the developed sensibility required to question, hear, and appreciate a wider set of conceptual possibilities for music itself. She has learned to perceive what without this initiation she would never have perceived. To progress to this level of mastery requires the prolonged, disciplined direction of attention and an imaginative expansion of their understanding of the possibilities of music. To fail to attend properly, to lapse in discipline, to shrink rather than expand one's imaginative capacity is to fall short of mastery. Seeing the progression of musical

skill this way can lend credence to Murdoch's idea of occupying *different worlds*. As the amateur becomes a master, they become, in some sense, a full-fledged participant in the world of music.

Yet even masters need to practice. The stepwise progression of learning the cello is a helpful example of how the imaginative work of loving attention is infinitely perfectible. Though the master player has attained a greater appreciation for the value of music than she had as a novice, there is no definite *end* to her study. Masters practice just as much, if not more than students; they do so not only to perfect their ability to reliably perform, but also because there is no limit to the possibilities for musical growth: their imaginative grasp on the meaning of *sorrow*, *joy* and *play* in music can always be strengthened. Appreciation of music highlights how understanding the value of certain activities requires both the disciplined direction of attention and the imaginative acquisition (strengthening and expansion) of the meaning of concepts.[127]

To initiate and then sustain this level of disciplined attention is to take a stance toward music akin to love. To love in Murdoch's sense is to engage in the work of imaginative understanding, striving toward an ever-deepening understanding of the possibilities of meaning in a particular activity, experience or relationship. This loving gaze is directed not toward any definite end point, but toward an ideal limit. The progressive attempt to bring the value of an activity like cello playing into view is infinitely perfectible because the set of possibilities for understanding music itself is infinite.

[127] The same process can be illustrated, Murdoch (1999b, 216) thinks, through the appreciation of human beings. While "we all have an indefinitely extended capacity to imagine the being of others […] there is no prefabricated harmony, and others are, to an extent we never cease discovering, different from ourselves." This is what she calls the "infinitely extensible work of imaginative understanding, of two irreducibly dissimilar individuals. Love is the imaginative recognition of, that is respect for, this otherness."

V

There is one final twist to Murdoch's account of practical thinking: the better our conceptual grasp on a particular concept, the more we understand its depth, context, history and unique value, the *less* intelligible these concepts become to the public. While contemplation and practical experience can enhance our expressive powers (e.g., Murdoch was also a gifted novelist), Murdoch claims that conceptual mastery is a fundamentally personal project.

Murdoch's view comes in response to what she calls the *genetic* argument about concepts. The genetic argument is this: consider the concept "green". The concept cannot be the name of something private. This is because the structure of the concept is just its public structure, which is established by coinciding procedures in public situations. On this view, Hume was wrong to worry about the missing shade of blue, not because a person could or couldn't picture it, or that we could or couldn't be persuaded that they had, but because the inner picture is *irrelevant*. According to the genetic argument, the possession of a concept is a public skill. What matters is whether I obey the traffic lights, and not my inner imagery or absence of it. On this view, I identify what my senses show me, placing a sensation under the concept "green" by means of the public schema which I have learned and internalized. Indeed, the very test for whether I have learned it or not is public (whether I respond appropriately to green in certain contexts).

The genetic argument seems rather intuitive. While surely we have private introspection, how can concept acquisition be private? After all, one obvious measure of one's conceptual grasp of a concept is their ability to *use* it effectively with others. But on Murdoch's view, to discover a particular truth is part and parcel of concept acquisition. If we have never met an honorable person, then we will forever miss something significant about what *honor* means. Moreover, coming to understand the *honorable*, is an inexhaustible task because the reality of its object, an

individual person in the world, is also an inexhaustible source for understanding. The requirement of practical experience with an object, activity or person in order to understand its significance looks familiar: Murdoch seems to be highlighting instances of categorical species, which we discussed in chapter one. To put oneself in a position to see this unique reality, Murdoch argues, requires first-personal engagement. But since such understanding requires more than words, but also loved experience, as we move forward into particular understanding we become less able to express the truth one sees, to move into an increasing conceptual *interiority*.

> Each of us lives and chooses within a partly private, partly fabricated world, and although any particular belief might be shown to be 'merely fantastic' it is false to suggest that we could, even in principle, 'purge' the world we confront of these personal elements. Nor is there any reason why we should. To be a human being is to know more than one can prove, to conceive of a reality which goes 'beyond the facts' in these familiar and natural ways. This activity is, moreover, usually and often inevitably, an activity of evaluation. (Murdoch (1999b), 199)

Murdoch's vision of practical thinking begins with the claim that all thinking is moral thinking. Such moral thinking, when it aims deliberately at truth, comprises the ongoing, imaginative work of loving attention. For Murdoch, puzzlingly, this imaginative work of coming to understand a concept makes us increasingly less able to express it.

To finish unraveling our initial tension, we need answer one final question: If the deepening of one's grasp on moral concepts is a discovery of what is there to understand and value in the world, and this discovery makes is *less* able to communicate the value that we are in

a position to see, why should we credit incommunicable values? If Murdoch is right, and the development of moral sensibility is at odds with our ability to articulate *what* we are sensitive to, then it seems difficult to see how and why we ought to credit such incommunicable values in the domain of moral justification. To answer, we must consider a case where the move into privacy—meaning inability to express the value one is in a position to see—is part and parcel of the move into full moral development. Consider her case of M and D.

> A mother, whom I shall call M, feels hostility to her daughter-in-law, whom I shall call D. M finds D quite a good-hearted girl, but while not exactly common yet certainly unpolished and lacking in dignity and refinement. D is inclined to be pert and familiar, insufficiently ceremonious, brusque, sometimes positively rude, always tiresomely juvenile. M does not like D's accent or the way D dresses. M feels that her son has married beneath him. Let us assume for purposes of the example that the mother, who is a very 'correct' person, behaves beautifully to the girl throughout, not allowing her real opinion to appear in any way. We might underline this aspect of the example by supposing that the young couple have emigrated or that D is now dead: the point being to ensure that whatever is in question as happening happens entirely in M's mind [...]

But as time passes, gradually, M's impression of D begins to change:

> [M] is an intelligent and well-intentioned person, capable of self-criticism, capable of giving careful and just attention to an object which confronts her. M tells herself: 'I am old fashioned and conventional. I may be prejudiced and narrow-minded. I may be

> snobbish. I am certainly jealous. Let me look again. Here I assume that M observes D or at least reflects deliberately about D, until gradually her vision of D alters. [...] D is discovered to be not vulgar but refreshingly simple, not undignified but spontaneous, not noisy but gay, not tiresomely juvenile but delightfully youthful, and so on. (Murdoch 1970, 17-18)

Taking M as our exemplar, what does a good practical thinker look like? At the very least, she is reflective, deeply connected to and critical of her evaluative emotions and capable of an inner progress toward a certain ideal way of seeing the world. But what kind of progress? If we conceive of M as wholly identifiable with her will and morally responsible only for her outward dealings with others, if she were "pure will", there is no progress to speak of. Since nothing in her outward actions has changed, she, in a sense, has not changed. But Murdoch's initial point here is that M has *obviously* changed in a deeply familiar and important way. Where she used to feel disdain and righteous contempt for D, M now feels attachment and regret at her inattentive judgments.

Where the moral philosopher struggles in vain to find some *action* to identify with M's moral progress, Murdoch points out that "M has in the interim been *active*, she has been *doing* something, something which we approve of, something which is somehow worth doing in itself. M has been morally active in the interim […]"[128] Murdoch tells us that M is "continually active, as making progress, or of her inner acts as belonging to her or forming part of a continuous fabric of being [...]."[129] For Murdoch, practical thinking is not always constituted by a series of

[128] *Ibid*., 19-20.
[129] *Ibid*., 22.

disconnected occasions for action. Instead, it can be a continual progression toward an ideal disposition to *see* the world as it really is. "Is not the metaphor of vision almost irresistibly suggested to anyone who, without philosophical prejudice, wishes to describe the situation? [...] M *looks* at D, she attends to D, she focuses her attention. M is engaged in an internal struggle."[130]

The imaginative work of loving attention is thus a "process of deepening or complicating, a process of learning, of progress, which may take place in moral concepts in the dimension which they possess in virtue of their relation to an ideal limit."[131] Moving from thought to action, on this view, is a progressive attempt to bring the world into view, which involves much more than deciding how to move one's limbs. Mark Hopwood marks the shift here nicely. For him, most moral theorists write that, in addition to doing the right action, "virtue simply consists in arriving at the same description of one's situation that an ideally rational agent would give." In contrast, Hopwood points out that for Murdoch, "the description itself is not what is most significant. Morality is less about the attempt to master a set of rational norms, and more about the attempt to do justice to the reality of other individuals."[132] I think Hopwood is right that for Murdoch, doing so involves coming to acquire moral concepts with which to understand what we are seeing, judging and doing. At this stage, the particularist might interject that what virtue consists in, in any particular case, depends quite a lot upon a holistic consideration of the individual deliberating as well as the situation that they find themselves in. Such progress is unlikely to be codified in terms of rules for conduct that can apply across situations.

But despite this similarity with the particularist, Murdoch's work on practical thinking marks a change in both the mechanics of how we encounter the world and decipher what is

[130] *Ibid*., 22.
[131] *Ibid*., 31.
[132] Hopwood (2018), 496.

significant in it and the aim and focus of our subsequent deliberation. For her, practical thinking is an ongoing activity, not a singular event, the aim of overcoming one's own prejudices, one that requires the agent to learn new concepts for describing the object of attention.

In contrast, mainstream moral theorizing involves a foundational commitment to a picture of practical thinking where ideally rational moral agents are subject to an independent, objective reality. For Murdoch, the reality to which we are subject, to which we are attending in life, is the reality of ourselves and other people, not merely a given world of objects. Murdoch sees the relevant attempt to capture reality as infinitely complex such that doing justice to it is an "endless task", one that does require a holistic evaluation that is best articulated by the particularist. But unlike the Dancy-style particularist, Murdoch consistently rejects the picture of practical thinking that grounds his view.

Instead, Murdoch takes her guidance and analogous comparisons from the world of art. The imaginative attention to the inexhaustible task of coming to know the reality of other human beings is a continual, creative act that is never finished, codifiable in rules or best understood as an attempt to capture a photograph realism.[133] The work of Murdoch's moral agent, like the artist, is to challenge and then reimagine the lenses through which they view the world, to imaginatively remake the world they see in light of the virtue they hope to embody.

With this analogy in hand, let us return one final time to the case of M and D. What is M attending *to*? What is she attempting to *do*? According to Murdoch, M is attempting

[133] Murdoch's conception of love is deeply influenced by the haunting insight of Simone Weil. For example, compare the former's talk of love as the inexhaustible work moving beyond ourselves, of attentive imagining an other to the Weil's version: "To empty ourselves [*Se vider*] of our false divinity, to deny ourselves, to give up being the center of the world in imagination, to discern that all points in the world are equally centers and that the true center is outside the world, this is to consent to the rule of mechanical necessity in matter and of free choice at the center of each soul. Such consent is love. The face of this love, which is turned toward thinking persons, is the love of our neighbor." (WFG 100)

> not just to see D accurately but to see her justly or lovingly. [...] Freedom is not the sudden jumping of the isolated will in and out of an impersonal logical complex [called moral deliberation], it is a function of the progressive attempt to see a particular object clearly. M's activity is essentially something progressive, something infinitely perfectible. [...] M is engaged in an endless task. (Murdoch 1970, 23)

Murdoch's point is that the significant and praiseworthy change in M cannot be captured by an account of practical thought whose structural focus is built around publicly observable actions that fall out of isolated instances of deliberation. Instead, she thinks that we are always adjusting our practical moral vision, that the practically salient features of ourselves are not captured solely by talk of reasons, and that the task of clarifying our vision of particular objects is not complete when conscious deliberation issues in overt action.

Recall the initial tension we identified in Murdoch's view: how is it possible that practical thought is both a (re)creation and discovery of the reality of the world? Let us now attempt an answer. We saw that Murdoch believes that all cognition is moral cognition. This means that the concepts we deploy in understanding the world are always already infused with moral sensibility. Murdoch thinks that this is a process in continual flux; the recreation of our moral concepts can move us closer or further from appreciating the value that exists in the world. Such a process, we saw, requires the disciplined direction of attention, whose gaze is reminiscent of love. Such a loving, attentive gaze, opens up the imagination, extending and expanding our understanding of the meaning of the moral concepts that are part and parcel of a developing sensitivity to the world. When practical thinking goes well, Murdoch thinks that the recreation of our moral

concepts moves us closer to discovery—gaining a firmer, deeper grip on the meaning of the concepts that help us understand and appreciate the value of objects and pursuits in the world. When it goes poorly, when our moral understanding becomes coarse, shallow and limited, we move away from discovery and into fantasy. Both sides of the bifurcation involve modifying the conceptual world within which we reason, but only one side (attention, fruitful imagination and love) leads us to discovery.

VI

I think Murdoch presents a compelling argument to reject the standard account of practical thinking. Murdoch's discussion points toward the idea that goodness is a form of practical knowledge: a refined and honest perception of what is the case, whose recognition requires a patient exploration and a just discernment of what confronts each of us when we encounter another human being. Maybe this discussion could issue in a different form of particularism, but it would have to be one that rejects the account of how thought makes itself practical that is assumed by extant particularisms.

Murdoch's philosophical work also suggests a limitation on moral theorizing that, if taken seriously, would alter the scope, depth and applicability of the subject. She tells us that "moral philosophy should remain at the level of the differences between forms of life that exist between us and not try to get behind them to a single form."[134] Moral philosophy need not be a neutral analysis of the world or nothing. Instead, it can be both analysis *and* imaginative exploration, which, according to Murdoch, implies that it is "discontinuous with the rest of philosophy"

[134] Murdoch (1997d), 97.

because it *must* be studied in the context of serious historical, sociological and psychological investigation.

If taken to heart, these lessons suggest that placing Murdoch somewhere on the spectrum of particularism ignores her potent and interesting account of how thought makes itself practical. Murdoch offers us a completely different way of thinking about practical thinking, one that frames a complicated relation between everyday life and moral philosophy. By rejecting the characterization of actions and reasons in favor of continual imaginative activity, Murdoch undermines more than the rigid limitations of theorizing or the dangers of generalization; she also rejects the entire conception of practical thinking that is assumed by both the Dancy-style particularist and generalist alike. To take seriously her claims about the role of love, attention and vision in practical thinking is to, in the same breath, sweep away the ground upon these forms of moral theorizing rests.

The second major contribution is a set of conceptual tools. Murdoch replaces individuated action with ongoing activity, an isolated will with an emotionally connected one, and marks the continual progression toward a more just perception of the world with the terms *love, attention* and *imagination*. She thus provides a framework for describing what is happening in the cases that have concerned us throughout, where the value of an activity, experience or relationship cannot be fairly represented in the language of reasons. In these cases, practical thought is best described as a *continual activity* rather than a singular event, one that is always ongoing. That activity consists of more than merely encountering neutrally describable, given features of the world. Instead, understanding what is there to be seen involves acquiring and then learning to deploy moral concepts to the case at hand. This is an iterative and sometimes highly individual process, one that requires first-hand attentive, careful engagement with the activity in

question. It is also clearly an exercise of the *imagination*, since what is there to be seen is partly a function of the concepts we bring with us to the situation, and partly a function of how we interpret the applicability of those concepts. This dedicated, attentive and imaginative work is encapsulated by what Murdoch calls *love.*

Murdoch also offers an alteration of our conception of moral reflection and disagreement. For her, moral reflection is a mode of engaging with the world that we have to *achieve*, not a capacity for freedom of choice that we have in virtue of being, say, human or rational. Moral disagreements run much deeper than the failure to grasp the reasons another person has, reasons that are neutrally describable and thus (in principle) available to all.

> Moreover, if morality is essentially connected with change and progress [...] we do not simply, through being rational and knowing ordinary language, 'know' the meaning of all necessary moral words. We have to learn the meaning; and since we are human historical individuals, the movement of understanding is onward into increasing privacy, in the direction of the ideal limit, and not back towards a genesis in the rulings of an impersonal public language. (Murdoch 1970, 29)

Seen in this light, moral disagreements are differences of *total vision*, not merely differences of choice between an array of considerations that count in favor: since practical thinking is part and parcel of coming to master moral concepts, moral disagreements are not only a matter of selecting different objects as morally salient out of the same (neutrally describable) world but because we *see* different worlds.[135] According to Murdoch, the two parties of a disagreement are

[135] Murdoch (1999d), 81-82.

actually living in different worlds. Such a disagreement runs deeper than any tussle over which considerations are genuine reasons in relation to a given world of facts.

This suggests yet another contribution: if we accept Murdoch's conception of practical thought, then the boundaries between practice and theory begin to blur. This blurring of the division between theoretical thought and action is a distinctive and overlooked feature of her work, one that raises doubts about the clarity and distinctness of the category of practical thinking itself. Recall that on a modern account of practical thinking, the performance of an action and the perception of reasons are independent elements. But, Murdoch claims,

> The world which we confront is not just a world of 'facts' but a world upon which our imagination has, at any given moment, already worked; and although such working may often be 'fantasy' and may constitute a barrier to our seeing 'what is really there', this is not necessarily so. Many of the beliefs which are relevant to action are unlike disciplined scientific or scholarly beliefs. They are beliefs in the genesis of which active imagination will play a part which is not necessarily sinister. The formulation of beliefs about other people often proceeds and must proceed imaginatively and under a direct pressure of will. We have to attend to people, we may have to have faith in them, and here justice and realism may demand the inhibition of certain pictures, the promotion of others. Each of us lives and chooses within a partly private, partly fabricated world, and although any particular belief might be shown to be 'merely fantastic' it is false to suggest that we could, even in principle, 'purge' the world we confront of these personal elements. Nor is there any reason why we should. (Murdoch 1999b, 199-200)

This view of practical thinking necessarily involves theoretical thought—the continual striving toward a reality that is itself already made over by our contemplative imagination when we begin to reflect upon it. As we saw, what we see as neutrally describable facts, include and are shaped by what values we are attracted to exist in a world we have already (partly) made. Whether we see another person as a lover or a stranger, a situation as an opportunity or an obstacle, is in part a product of our mastery of the moral concepts we deploy as descriptions of what is before us.

This view raises an important question once again: if Murdoch is right and an increasing cognitive grasp on a moral concept leaves us increasingly less able to express the value we see, why should we credit such values in interpersonal moral justification that only we are in a position to see? One answer lies in reflection on the nature of intimate relationships. In chapter one, we saw that the value of an individual human being is only fully available for appreciation from within the first-personal attempt to bring that value unto view. Absent any practical acquaintance with the relevant person themselves, it is difficult to see how we can be said to know them. Since what we value from within that relation is not always the generic propositions that we can articulate in words, we concluded that the inability to articulate the value of an individual human being is no obstacle to appreciating that value.

Applying this example to Murdoch's view: it would be a mistake to say that since an agent is unable to articulate the value they see in their beloved that we should give less credence to their claim to see and appreciate that value. Instead, following Murdoch, part of understanding the depth and character of the value of an individual human being is an *achievement*, something that takes attention, discipline and an imaginative openness to extending the possibilities of our moral conceptual understanding, then it seems plausible that those who have achieved (or come closer to achieving) such an appreciation will be unable to share that standpoint in publicly

intelligible words. If the value in question cannot be simply *given* but must rather be *sought*, then we have good reason to hold that there are important, existent and familiar values that cannot be expressed in terms of publicly appreciable reasons for action.

The question above assumes that credence in the significance of a value should move inversely proportional to an agent's ability to express it in terms that any rational agent could be expected to understand. But to do so is to fail to see Murdoch's vision of practical thinking, one that is not limited to the apprehension and articulation of reasons. Hers is a unified, continual activity whose aim is the progressive clarification of the reality of individuals in the world, whose features are there to be appreciated, but only by the imaginative, the attentive, the clear-sighted.

Murdoch's view of practical thinking is important, in part, because it shows us, in vivid detail, that there are alternative ways of understanding morality in human life, beyond what has gripped contemporary ethics for the better part of the last two centuries. Beyond, that is, outlining the contours of our duties, naming what we owe to each other. Murdoch's work shows us how ethics can also be about striving to live a conscientious life in every respect. Her view is aspirational, aiming at the infinite perfectibility of our comportment with all our fellow human beings. Reviving this ambition is part of the contest of this dissertation. If philosophers stick with the prevailing theory of the aims and boundaries of the dominant strain of ethical theory, which seek to outline principled standards for how we ought to comport ourselves with strangers, it will be difficult to prove conclusively that these demands of morality always trump other considerations.[136] The Murdochean idea of the pervasiveness, inexhaustibility, and perfectibility of moral concept acquisition in practical thinking represents a more plausible set of demands that could systematically trump other considerations.

[136] I am by no means the first to make this point. See especially MacIntyre (1981; 2018) and Williams (1985).

We now have the conceptual tools to move forward. Recall that our initial diagnosis in chapter one located the problem of artificially restricting what counts as a genuine reason within a set of methodological issues with the way mainstream moral theorizing represents reasons. We then turned to a possible solution with particularism, but ran into another roadblock: particularism's most prominent view seemed to subscribe to an idea of practical thinking that obscures the very phenomena that sparked our investigation. This chapter's aim was to lay the groundwork for a new methodology for representing certain instances of practical thinking, one that could fit nicely with some form of particularism to better account for certain central values in moral life. Murdoch's account of the activity of love, attention and imagination presents some enticing tools for capturing modes of valuing that concern us.

In the next chapter, I put these tools to work. With them, I create a fresh account of practical thinking that moves beyond talk of reasons to amend the reigning philosophical methodology for building theories. In doing so, I'll show how what I call *loving attention* can amend the traditional view to better describe what is going on when we study philosophy, appreciate art or love one another. The goal will be to better capture the mode of valuation that grounds the example cases we highlighted in chapter one.

Chapter 4: Practical Thinking as Loving Attention

I

This dissertation began by highlighting a set of familiar activities, experiences and relationships whose true significance escapes meaningful representation in terms of reasons. The generality of this problem revealed its ground: a shared methodology among contemporary analytic moral philosophers, consisting of two ways of delineating reasons. Investigating what I called the formal and substantive approaches to representing reasons unearthed a shared conception of practical thinking whose central focus is reasons for action.

My conclusion was that this picture of reasons in practical thinking works to minimize certain ways of valuing that are centrally important in our everyday lives. I suggested that what philosophers take to be fit for serving the role of reasons itself embodies a commitment to some unacknowledged substantive claims. In order to serve as a real reason, a consideration must be like an ordinary object: in some sense there (be it given or constructed) for all competent, rational users of descriptive language to see. A reason, so construed, functions like an interpersonal commodity, producible to anyone on command; morality is correspondingly taken to be an marketplace of ideas that encompasses all the world with reasons as its currency. Since this moral marketplace includes all of us, in order to meet the fee of justification, the moral agent must be able to offer up the requisite currency; that is, to articulate one's reasons such that they meet this interpersonal standard.

Where the formal and substantive approaches assume that talk of reasons rests on some fixed, stable features available for all to see, the cases we considered suggest otherwise. The values that the activities we encountered attempt to bring into view are not always articulable,

not simply given by the world, but are rather revealed from a fully engaged perspective. Since the value of such pursuits like loving relationships, the appreciation of art or philosophy is yet to be fully revealed, there can be no shared understanding of the full value of the activities absent the first-personal attempt to dialectically engage with them, to bring their unique objects into view. If philosophical models of practical thought are to avoid distorting or passing over the value of these centrally important activities like studying philosophy, appreciating art or loving individuals and thus improperly weighting the reasons that flow from these values, theorists need a way of representing why we act in light of activities whose value we cannot yet appreciate or explain. Specifically, theorists need a way of capturing the value of these pursuits that does not rely on any of three assumptions outlined in the previous chapters: that legitimate considerations (i) must rest on some deeper neutrally describable features of the world, be (ii) articulable in advance of some particular practice experience or (iii) rely for their legitimacy on a robust, shared understanding of the good.

This chapter presents an alternative view of practical thinking that attempts to move beyond these assumptions above, in fact, beyond talk of actions done for reasons entirely. I call this alternative account *loving attention*. On its own, loving attention provides a richer picture of what is happening when we come to value that which we do not yet understand. It also represents a more compelling argument against the use of general moral principles than the most prominent particularist account we considered in chapter two. Finally, as a model of practical thought, this account provides the resources for a richer understanding of the role of moral concept acquisition, imagination and attention.

II

We all face moments in life where reflection prompts a total rethink. More than means-end deliberation, this type of deliberation is a revisioning of what is seen when one looks upon the world; change of this sort marks an upheaval in the brute orienting concepts one deploys to describe *what* is there. Such a *paradigm* shift involves questioning those deeply internalized conceptions of the good that structures one's perception of the world. For some, these are moments of peril when the conceptual ground one walks upon begins to quake—what appeared to be a stable relationship turns out to be a castle made of sand; what was once evidence of divine intervention is now ordinary cause and effect. For others, these can be moments of thrilling discovery—what was once a meaningless rectangle on the wall becomes an aesthetic object worthy of lifelong study; a series of chance encounters with another person sparks a relationship once thought impossible; a stranger becomes our beloved. These are alterations in more than our reasons—they are changes in our grasp of the good in life, changes in *that from which our reasons flow*.[137] Such moments mark a sea change in what life itself is like, how we see each other and our understanding of the good.

When deciding what to do in such cases, where the concepts that order and structure our understanding of the good are challenged, talk of *reasons for action* can seem out of place. This is because deliberation of this sort is not neatly contained within the logical space of practical reason. Reasons are called down from concepts one already sees as good, the ones that *already*

[137] Those who argue for what is referred to as "reasons fundamentalism" would take issue with this claim. A reasons fundamentalist holds that it is reasons, not values or the good, that are the fundamental unit of analysis in moral theory. Scanlon (1998; 2014) holds this view. Some have thought reasons fundamentalism to be advantageous for building workable theories like contractualism, since it can ground moral *agreement* absent any robust agreement about the human good. As a result, reasons fundamentalism can claim a plausible neutrality with respect to the good life. See Hieronymi (2011). Others emphasize a corresponding disadvantage: accounts that put reasons first: they lend themselves to an unsatisfying *thinness*. Fernandez (2018) argues that since reasons are responses to what we see in the world, right action should be derivative from understandings of the good or the good will in theoretical accounts of moral deliberation, not the other way around.

orient and describe our sense of the good in life. But in moments of radical doubt, conversion or discovery, these are the very grounding concepts that are under review. New kinds of goods are being discovered. There is, then, no reason (grounded in some deeper, previously settled evaluative framework) to point to in virtue of which one acts, because what one wants, and indeed the relevant settled evaluative framework itself from which one typically acts in virtue of reasons, is precisely what is being called into question. This also means that the activities, experiences and relationships which prompt this way of reorienting are not valuable *in virtue* of some deeper, neutrally describable features that hold still through time. To see them as such is to import an already settled classification of the kinds of meaning available for an evaluative term to have—this will not do.

Such paradigm shifts, those moments that open up new conceptual possibilities often occur after a period of progressive, dialectical engagement in the activity itself. The activities we considered in chapter one have an essentially dialectical structure such that full appreciation of the good(s) they offer is only attained stepwise through a continual process of attentive investigation. In such cases, the process of reorienting ourselves in the world, how and what one sees in it, cannot be specified in advance of the attempt to bring it into view.[138] It appears, then, that in these moments of deep discovery, *reasons talk runs out*. Following suit, the philosophical methodology that models practical thought as essentially actions done for reasons cannot easily capture what is going on.

---

[138] Callard (2018) names this process of seeking out some value or set of values that we do not yet understand as *aspiration*. But where Callard thinks, in order for an act to count as an instance of practical agency, it must be responding to some *reason*. My view holds that some of the phenomena she considers (love, appreciation of art, philosophy, skill acquisition) are not best captured by talk of reasons at all. This is because the language of practical rationality brings with it all the conceptual baggage outlined in section one. I think we do better to hold that the agent is responding to the stepwise revelation of some previously undiscovered good, orienting themselves toward it through a continual process of loving attention. Callard and I agree, however, that the end result of this process of coming to value that which we do not yet understand, whatever we call it, can be reasons.

The view I discuss in the following section looks different. Loving attention is a progressive attempt to appreciate the significance of an activity; it represents an altogether different mode of valuing than coming to understand the force of reasons that are grounded in given features of one's circumstances. In the latter case, it is assumed that the circumstances one finds oneself in unproblematically yield features that ground the application of the term *reason*. Given this framework for reasons, this means that the value or good of the given pursuit is somehow *there* for any rational agent to see (if given an adequate description that includes the relevant psychological states of the agent) since the features that ground the evaluative designation of a consideration as a reason are *given* by simple appraisal of the world. However, as we have seen, there are certain values that are not *simply there* for all to see.[139] To see things this way is to assume that there exists some set of concepts, which always sit beneath the evaluative good the practical thinker sees that fundamentally structure practical thought.

As we have seen, this is not always the case. This picture of practical thought is incomplete. One reason why is that it is saddled with an inadequate philosophical view of language, which ties action description to a narrow range of kinds of classification. For instance, to understand the concept *human being* is not merely to grasp a biological taxonomic classification system *plus* an evaluative extra. It is instead, as Cora Diamond puts it, "to know how thoughts and deeds and happenings, and how happenings are met, give shape to a human story; it is a knowledge of possibilities, their weight and their mysteriousness." Importantly,

[139] The phrase "out there" may be misleading. While it may appear that my argument here is really an objection to realism about values, I think it applies equally well to the noncognitivist. While the realist and the noncognitivist disagree about the origin of value (are they in the world or in us?) that difference makes no difference to my argument, since I am describing the phenomenal experience of appreciating value. Whether values are, in fact, given by neutrally describable features of the world or invented by some noncognitive feature in us, both camps agree that we typically *feel* as though the demands of moral life are discovered in the world. See McNaughton (1988), especially chapter 6, for a discussion of the call and response between the realist and noncognitivist.

neither is the concept of a human being simply *homo sapien* "plus an evaluative extra."[140] Indeed, to grasp a biological concept is *nothing* like what it is to grasp the human concept. To map this distinction onto our discussion from chapter one, Diamond appears to be arguing that what it is to acquire the concept *human being* is not an instance of *accidental* but rather *categorial species*. The non-identification of mere naming and concept acquisition is illustrative:

> grasping a concept (even one like that of a human being, which is a descriptive concept if any are) is not a matter just of knowing how to group things under that concept; it is being able to participate in life-with-the-concept. What kinds of descriptive concepts there are is a matter of the different shapes life-with-a-concept can have. Life with the concept human being is very different from life with the concept member of the species Homo sapiens. To be able to use the concept 'human being' is to be able to think about human life and what happens in it; it is not to be able to pick human beings out from other things or recommend that certain things be done to them or by them. The criticism I am making could be put this way: linguistic philosophers have brought to their study of language an impoverished view of what can be involved in conceptual life. (Diamond 1988, 266)

I think a similar connection holds between coming to value some object, experience, activity or person and acquiring the requisite concepts with which to articulate and properly see them. Neither are processes by which the mere naming of a category can bring an understanding of its practice into view. Love, philosophy, art, or beauty, like *human being*, are not terms that can be

[140] Diamond (1988), 265-266.

applied to a piece of language *in virtue of* the truth or falsity of premises or the validity of inferences. Instead, these concepts, like *human being*, are built up from one's sense of how one lives, that is, what life is like. Philosophy encourages the idea that such practices and judgments *have no part* in philosophical thinking, that the practice of moral theorizing does not call on a responsiveness involving our whole selves, but rather only to ourselves as rational agents. The approaches to the standard view of reasons we have considered, together with the corresponding picture of practical thinking that sits in the background of the dominant methodologies for discerning them, is representative of this general attitude in the field. The range of types of concepts on offer is thus limited by the view of language generally affirmed by contemporary analytical philosophy and does not accommodate the breadth and possibility of actual language use. This leads philosophers to take great care to define terms like *reason*, *obligation*, *rationality* because, so it is thought, the sparser our philosophical ingredients, the better for theory. But if we step outside the realm of philosophical debate, and look plainly at the phenomenon moral philosophy is allegedly about, we might ask: *why would anyone think less is more*? If, following Diamond,

> we get rid of the idea that using a concept is a matter of using it to pick out what falls under the concept and what does not, if we see instead that life with a concept involves doings and thinkings and understandings of many sorts, into which one's grasp of the concept enters in different ways, then we can accept that coming to understand a conceptual life other than our own involves exercise of concepts belonging to that life. (Diamond 1988, 276)

I think the deployment of reasons talk is an instance of this larger phenomenon Diamond attempts to bring into view. Recall our discussion of *deflection* in the introduction: Diamond's point here is at once a philosophical and a cultural one, that moral philosophers suffer "from a form of cultural deprivation rooted in philosophical practices. [Their] case is not a case of as-if loss of concepts but a particular kind of actual loss or failure of mastery."[141] In my examples, the exclusive reliance on the concept of a reason to describe substantially richer experiences of coming to value new activities, experiences or persons represents a particularly sneaky inarticulateness in modern moral philosophy, one that has become so ubiquitous that it is very difficult for us to get in view. I have made the argument (in chapter one) that using *reasons* as the sole basis for interpersonal justification represents a general loss, or perhaps a *stretching thin* of a concept to cover too wide an area. Given only the language of reasons, some values are, as Diamond writes, "as far as modern ethics is concerned, no longer capable of expression; the moral experience of ordinary people cannot be illuminated by a vocabulary used in the ways available to modern philosophy." This view, whose inarticulateness about many complexities in life, presents a cultural ideal "as a philosophical necessity, [which reflects] that picture of the human personality which our culture in general has inherited from the Enlightenment."[142] In the cases we have been considering, talk of actions done for articulable, publicly comprehensible

---

[141] *Ibid*., 261.

[142] Murdoch's views on the scope of this problem are unclear. In *Sovereignty of Good* (1970) she frames her criticism as leveling an *amnesia* amongst largely atheist, Oxford educated philosophers such that they simply *write as if they forgot* the Christian historical cultural tradition in which their ideas of *obligation* and *authority* properly fit. In other places, (1999e) however, she traces the philosophical difficulty to a deeper problem in Western culture. Diamond summarizes Murdoch's criticism this way: "We see ourselves as capable, using reason and our senses, of describing the world; and we take the application of scientific language as the model for such description of reality. We contrast with that side of ourselves (reason and perceptual capacities) emotion, desire and will, and we explain to ourselves all the rest of our uses of language as involving those latter parts of the soul. In the moral life of beings conceived in such a way, there is no need for moral concepts other than the most general ones like good and right, together with straightforwardly applicable descriptive concepts. Philosophers, then, who accept such a picture of the soul will be blind to uses of words for which that picture leaves no room. [...]" (Diamond (1988) 262-263)

reasons is strained at best and obfuscatory at worst. In its place, what I will call the imaginative work of loving attention can stand as an enriching amendment to the picture of practical thinking that sits in the background of much of contemporary normative moral theorizing. Let us sketch this alternative.

III

Practical thinking begins well before the adoption of an intention to act. Before one can form an intention, one must first have some sense of what is the case. This requires acquiring certain concepts and learning to properly apply them. The acquisition and deployment of concepts is, among other things, the development of a *way of life*. This means developing a situated perspective from within a set of cultural practices; it is a matter of building up, understanding and then learning to deploy the most general of concepts like *right* and *good* in the context of daily life. It also requires a wide range of more specific, secondary concepts depending on the situation, like *trust*, *sincerity* or *honesty*. Doing so, in some cases, requires just that: *doing*. Consider once again Diamond's view of how we tighten our cognitive grip on the concept *human being*. An essential part of deepening one's grasp on the concept of a *human being*, we saw, necessarily requires some practical experience. To understand what a human being is involves more than a grammatical grasp on a taxonomic classification system plus the relevant differentia for the species *homo sapien*. Instead, understanding in this sense is about how agents relate to real thoughts and deeds and coincidental happenings, and how happenings are met, give shape to a human story. To deepen one's understanding of a human being is to gain insight into the wide possibilities of human life. Learning a way of life—familial and friendship relations, seeing *this* as dangerous, *that* as enticing—is only possible by moving through the medium of activities,

experiences and human relationships. We must live with each other to learn what the world is like—this is precisely how we describe and evaluate it. Acquiring certain concepts is a matter of engaging in practical thought in the process of learning a way of life.

Learning in what sense? Learning involves, among many other things, extension beyond what is already known. This is an essentially imaginative process that is both *receptive* and *creative*. Sensitive receptivity to the minute details of the world, internalizing what is *salient* in each case, to the best of one's ability, is surely important. But as we saw in the previous chapter, bringing the world more clearly into view is also a matter of actively reframing one's understanding of what is there to see. Such reframing is the difference between a child who is angered by their parents' rigid rules and the adult who looks back, after years of experience and feels gratitude for the inculcation of strong habits of mind. Doing so requires taking a certain *stance* toward the object of practical thought, one that is attentive and imaginative.[143] In the cases we have been considering, the stance required to sustain the wholehearted attempt to bring the world more clearly into view through attentive, imaginative effort and a process of continual receptive reflection can be described as *loving attention*.

What does loving attention entail? Consider an example from the world of art. If, when I gaze upon a Rothko, *all I see are two rectangles*, then I cannot attend to its artistic features; I do not even know that they exist. This is because I do not yet have the conceptual resources to appreciate what makes such a work valuable. But where to begin? Rectangles of color, seen as just that, are simply not worth attending to; coming to see them as more than that requires

[143] Bremner (2022) provides a unique approach to thinking about moral salience, showing how Kant's framework of reflective judgment provides for the revision of principles that brings in the insights of particularist understandings of moral salience. This means understanding moral deliberation as "both the incorporation of new moral particulars and the revision of moral universals on the basis of such discoveries." While I am sympathetic to this hybrid approach to understanding moral salience, I am skeptical that the framework of Kantian reflective judgment really "incorporates new particulars" in the sense desired by particularists.

developing a way of looking, paired with building up the language of artistic expression: one must do a lot of looking *and* thinking to see with new eyes.

What is there to see? One can learn to look for color, contrast, composition, line, scale, materials, context as independent yet interactive tools in the artistic workshop. In and through this practice of looking one also learns to *speak* of them, to acquire them as meaningful concepts. Where do such concepts come from? There are two different, yet related sources. The first and most obvious is through conversation and social habituation. In the case of art appreciation, this might take the form of an art class, a series of field trips to museums and a whole set of conversations with art lovers. Through this rich social context we can gain insight into what one ought to look for in particular works of art.

But this is only a start. While it is surely possible to glean a generic set of strategies, artistic qualities and useful tips from our social surroundings, more is required in order to appreciate a work of art for oneself. As we saw in the previous chapter, the imaginative, recognitional task of value apprehension is carried out through the first-personal, disciplined direction of attention. The specific features one attends to are limited by the acquisition and contextual placement of (moral) concepts one has at the time. In other words, one can only attend to the conceptual world one can see.

Importantly, this conceptual building up does not fit neatly into a single episode. While one might experience an "aha!" moment, where all the previously muddy concepts *click* together, such moments mark a waypoint, not the beginning, of value appreciation. Seen as the continual, progressive process of attending, My vision of practical thinking in these cases, following Murdoch, is not a time-bounded event, but rather goes on all the time; it is a piecemeal business of building up structures of value through our habitual conversations and relations with others,

not something that is switched on in a moment of deliberation. More than an occasional, discrete task to be done, the activity of loving attention is an infinite and inexhaustible task of attending to the world as it really is.[144]

What is being attended *to*? On the surface, it might seem as though one can attend to just about anything. In the case of artistic appreciation, we can point to features of the work, its creation, context, the intent of the author, etc. Are there limits on what can be an object of appreciation in everyday life? In the cases I have been discussing, the object is less clear than the case of a painting. After all, when one begins down the path of discovery, the object of attention is not yet fully in view, and indeed this inability to fully appreciate the object of attention makes it all the more difficult to state what exactly one is attending to. As a preliminary start, we might say that when one begins down the path of appreciation for art, philosophy or another human being, they are attempting to attend not merely to the surface level features of the object, but rather, as Gavin Lawrence states it, what *runs deep in human life*. On his view, they seek "a fuller description of lives, of the wider settings in which work-activities take place, their personal and social resonance, their historical valence, their richness or lack of it: in short, an understanding of the contextualism of value, in all its local, historical, detailed [...] aspects."[145] To achieve this widening of conceptual, historical, social and personal resonance is, in part, an exercise of the

[144] The language of "infinite" and "inexhaustible" sources of value is Murdoch's. While she is less than clear about what *she* means, I think the particularist is best positioned to make sense of this claim. After all, if one subscribes to holism about value and denies the codifiability of morally salient that flows from sensitivity to those complex and situationally-contingent values, then the idea of, say, coming to understand another, ever-changing human being in their entirety looks *practically* infinite.

[145] Lawrence (2018). Lawrence's claim here extends beyond the epistemological battleground between contextualism and forms of foundationalism. I think Lawrence is attempting to make a much broader claim about the incredible complexity and subtlety of moral life: a case for a contextualist version of holism about value that is equally at home with particularists, though not of Dancy's sort. He writes, "the meaningfulness of my activity takes on its valence as serious endeavor, through others pursuing theirs, as equally it can be cheapened and threatened ('the social medium'): a fluid yet partially structured holism. Such structure thus goes hand in hand with the contextualism of value, where the value, the goodness or badness, of things is a function of their role, their contribution, good or bad, in their context." (2018, 240)

imagination, which on my view has two facets. First, imagination involves coming to understand the worth of an activity, experience or relationship that one has not yet been initiated into involves first imagining the possibilities for what might lie ahead. While surely such imagining is often mistaken, narrow or prejudiced, it is this ability to peer into the unknown, to take what for now has little value or meaning and entertain the possibility that something hitherto unknown might emerge, that marks the starting point for embarking on such world-altering pursuits at all. Absent any imaginative gaze, there is often only an indefinite inkling to motivate the first lurching steps down a road whose destination is unknown.[146]

The second imaginative aspect of this activity involves our stance toward this dialectical process itself. This is the stance of openness to discovery. To continually challenge our known conceptual framework without guarantee of what lies ahead could be said to involve an imaginative faith.[147] Imagination because coming to see something as valuable for the first time involves imagining new purposes, ideas, ends and ways of understanding what is before you. Faith because what reason there is to pursue the activity in question is as yet unresolved—a dialectical process of understanding yields no guarantee to the beginner that there *will* be some payoff in the end. Since what value there is to be appreciated in the cases we have

---

[146] We might helpfully mark its significance by thinking of what it is to *lose* a concept. Diamond (1988) highlights a slew of ways in which we can have concepts without meaning, meanings without concepts or be at a loss for concepts to describe our experience at all. At the heart of each type of loss, Diamond tells us, lies a complex string of cultural changes that erase, diffuse or otherwise disrupt an individual's ability to understand their own experience. In the case of moral theory, I have been outlining a sort of supplementary version of Diamond's first case: philosophers who deploy the word "reason" to cover a whole swath of human life and experience, but define it in such a way that the concept is not up to the task. The question of how and to what extent the concept of a reason has been stretched, stuffed, gutted or otherwise morphed I leave unanswered for now. Just as concept loss involves a winding, historically situated set of social and cultural phenomena, so too does concept acquisition and shift.

[147] I point not to any religious faith, but to a confident stance of a scientist toward conceptual discovery that Berlin (1978) speaks of: "…that the facts to be fitted into the scientific grid and subsumed under the adopted laws or model (even if public criteria for selecting what is important, relevant, etc. from what is trivial, peripheral, etc. can be found and employed) are too many, too minute, too fleeting, too blurred at the edges. They crisscross and penetrate each other at many levels simultaneously, and the attempt to prise them apart, as it were, and pin them down, and classify them, and fit them into their specific compartments turns out to be impracticable."

discussed—their full form, content and relevance—is sometimes only available long after getting underway, the agent must take heart that *something more* will appear after becoming initiated into the relevant practice.

What makes this imaginative work of attention *loving*? We can begin by marking love as more than a positive emotional valence. While surely love involves a deeply held emotional connection, such affection can easily come apart from the dedication, devotion and other attitudes highlighted by the cases we have been considering. One can be devoted without joy, or feel attached without long-term commitment. The idea of *love* I am after also involves a settled disposition. Preliminarily, then, let us say that to be *loving* is to engage in a mode of interaction that is marked by settled disposition of care.

But this cannot be the end of the story. If being loving is a settled disposition, it might seem natural to call it a *perceptual* capacity. A critic might dismiss loving attention as a matter of perception at all. Consider Gilbert Harmon's rather revolting case: imagine you observe a group of people burning cats. On Harmon's view, one need not actually need to assume that burning cats *is* wrong in order to observe that it was wrong for those folks to harm the cat. On his view,

> observation plays a role in science that it does not seem to play in ethics. The difference is that you need to make assumptions about certain physical facts to explain the occurrence of the observations that support a scientific theory, but you do not seem to need to make assumptions about any moral facts to explain the occurrence of the so-called moral observations [...]. In the moral case, it would seem that you need only make assumptions about the psychology or moral sensibility of the person making the moral observation. In the scientific case, theory is tested against the world. (Harman 1977, 6)

Harmon is making a claim about the structure of moral thought. What to make of this case? Such a person, he might say, is not necessarily missing any *perceptual* capacity. After all, they are fully informed about what they *ought* to pay attention to, what is there for them to see.[148] In this case, they might believe the truth of the observation "it was wrong to burn those cats" but fail to believe that that proposition corresponds to anything in the world.

John McDowell would deny Harmon's analysis. For McDowell, the structure of moral thought does involve a perceptual capacity. In what sense has the person who fails to think that the wrongness of burning cats corresponds to anything in the world actually *perceived* the relevant object of attention? To borrow Murdoch's terminology: it seems that Harmon's agent is capable of *looking* but not *attending*. He gets the truth value of the moral proposition correct, but fails to see the corresponding value that is being violated in the world. In this case, the sense in which the obstinate unfeeling agent is *informed* is merely skin deep. For, if they truly were, then they would not be able to reject the reality of what is there to be seen, in this case the grotesque burning of animals. In response, we might distinguish a robust *perceptual capacity* from mere *acknowledgement*. For instance, following McDowell, we might say that what marks a virtuous person from a non-virtuous person is a difference in their perception of the situation, but in a deeper sense than mere acknowledgement.

> A kind person has a reliable sensitivity to a certain sort of requirement that situations impose on behavior. The deliverances of a reliable sensitivity are cases of knowledge; and

[148] Seitya (2013) lays out how we might account for this difference in a Platonic view of concept acquisition by distinguishing between the *ordinary* sense of understanding of a concept and *true* understanding. Since every concept brings with it the norms governing its use, on this view, we might say the informed xenophobe has an *ordinary* but not fully *true* sense of their situation.

> there are idioms according to which the sensitivity itself can appropriately be described as knowledge: a kind person knows what it is like to be confronted with a requirement of kindness. The sensitivity is, we might say, a sort of perceptual capacity. (McDowell 1979, 331-332)

On this view, what it means to exercise a reliable sensitivity to the world is to have knowledge of what is called for in one's circumstances. While I suggested that love *does* involve some form of moral vision, a merely functioning perceptual capacity is not enough to garner the dialectical phenomena in question. Let us then add to our definition and say that love involves both a perceptual capacity in McDowell's robust sense, taken to be an openness to curious, beneficent and other-directed discovery combined with a humble recognition of the opacity of others.[149]

Beyond reliable sensitivity to the morally relevant details of the world, loving seems to require some recognition or affirmation of the value of what is being attended to. Velleman offers a particular account of the value of objects that love brings: "if love is a matter of 'really looking', then there must be something that the lover is looking at, and the most plausible candidate for this role is that which also functions as the object of respect, i.e., the rational will." For Velleman, love is the proper response to recognition of the rational will in others. But this does not seem quite right, either. On Velleman's Kantian view, love and respect are hard to distinguish. In fact, he thinks it is quite revealing that it takes a philosopher to see that a value we tend to have a perfectly good grip on in everyday life like, say, human dignity, is the value that *really* moves us to love. Since the object of each is the rational will, we are left with the question

[149] There are many more thorough and wide-ranging accounts of love than what we can discuss here. See Ebels-Duggan (2019), Hopwood (2018) and Silverman (2019). My focus here is on sketching just the framework of practical thinking necessary to make sense of the phenomena of love, not its full depth and content.

of how to distinguish love for a particular person from respect for the rational will in them. For Velleman, the distinction runs thus: where respect for the rational will in others acts as a check on our motives, limiting what we can permissibly do, love "disarms" our emotions, making us vulnerable to each other.[150] While one might surely agree that love *can* be disarming and that it surely *can* share some similarities to respect, I think that Velleman is wrong to see the two as sharing the self-same object. Where Kant takes respect for the rational will to be a standing universal demand, I see love as precisely *not* a universal, impersonal demand but rather a particular recognition of the unfolding value of an individual object.[151] Such a task is likely to be difficult and tedious, piecemeal and often require much time and effort to achieve. This is because coming to be in a position to see the value of an object is never as easy as "really looking" nor as simple as "acknowledging" some feature in another like the rational will.

While perception and recognition are part of the story, my suggestion is that loving is also a fundamentally particular activity. This means not only that its demands will vary in content and form so drastically from object to object, but also that its value is only appreciable from within that unique relation. I agree with McDowell that love, like virtue, must have a deeply rooted perceptual component. I also agree with Velleman that it must have a cognitive component. But to these I might add particularism's holistic, anti-codifiable frame, insisting that love can be a response *to the beloved themselves*, which is difficult to formalize, universalize or otherwise make suitable for a generic relation without losing sight of the object and value in question. I

[150] Velleman (1999), 94-95. Here Velleman is speaking of loving a human being. But we are interested in a much wider, though related sense of love. While we do not love philosophy as we love our children, the sense of love I am after the overlap between these various forms of love, that is, what if anything, they share.
[151] Hopwood's excellent (2018) discussion of Murdoch's line about love as the on the "stunning realization that something other than oneself is real" unpacks this term in much more detail than we do here.

think Velleman's view misses what is frequently the object of love, its grounded particularity in the dialectically unfolding phenomenal experience of loving an individual.

We now have a sketch of the broad contours of an account of practical thinking, which kicks in where talk of reasons run out. In such cases, where articulable considerations that capture the full force and form of a particular value are unavailable or impossible absent some direct, first-personal practical experience, we do better by speaking of the dialectical activity of loving attention. Such activity begins with the questioning of the brute-ordering concepts from which reasons (as *in-virtue-of* concepts) are called down. This deeper level deliberation is itself a process of imaginative discovery with passive and active components. Imagination involves a developing sensitivity, a coming alive to bits of reality previously hidden. It also involves a creative activity of ever-expanding categories and possibilities for human life, a gradual building up, through practical experience, of new moral concepts. This imaginative process requires practical experience because both its aspects do—that is, both the expansion and development of a reliable sensitivity to new details of the world and the active creation of new categories and possibilities under which to understand and deploy them—require lived practical experience. In this sense, *doing runs ahead of our concepts*.[152]

This imaginative process requires a multifaceted dispositional stance I called *love*, whose facets include an (i) an open attentiveness to the object, (ii) a development of a perceptual capacity to see the object of appreciation for what it is and what it could become, and a (iii) cognitive capacity to recognize and deepen one's understanding of the object of attention and

[152] Murdoch (1970), 43-44 makes this point beautifully: "I have suggested that we have to accept a darker, less fully conscious, less steadily rational image of the dynamics of human personality. With this dark entity behind us we may decide sometimes to act abstractly by rule, to ignore vision and the compulsive energy derived from it [...] But if we do leap ahead of what we know we still have to try and catch up. We cannot run very far ahead of knowledge, and attention is our daily bread."

therefore the proper treatment of it.[153] Call this iterative process of discovering value the imaginative work of *loving attention*.

IV

Thus far, we have established that there are certain activities, experiences and relationships whose value is only revealed via a dialectical process of attentive, imaginative engagement whose unifying attitude is reminiscent of love. This means that coming to see the value internal to these pursuits rests on our ability to continually attend via a process of deepening understanding and imagine that there is some internal good to be found in the pursuit, even if one cannot yet understand it. Done well, this process involves a sensitive, generous and loving gaze.

Let us return to our central cases with fresh eyes. Why study philosophy? Presumably, the answer is that the agent sees something worth going for. Put another way, they see some *good*. Since that good is not yet fully revealed, there is not, or may not be, some definite, articulable proposition that the agent can point to as a corresponding reason in virtue of which they act. The content of that good is only partially revealed to the philosophical novice. While they might be compelled by its provocative questions, moved by the logical structure of fine argument, or may be thrillingly challenged by some counterintuitive conclusions, the novice lacks the conceptual tools to understand the full-fledged goods on offer. Learning to read carefully, deconstruct

[153] I am thinking here of Talbot Brewer's reflections on Ann Carson's work on *Eros* as a *stereoscopic* vision. For Carson, stereoscopic vision is an experience of looking at one's beloved that reveals "[t]he difference between what is and what could be is visible. The ideal is projected on a screen of the actual, in a kind of stereoscopy." (Carson (2014, 17). Brewer (2022, 4) builds upon this notion by outlining stereoscopy as a mode of contemplation that reflects Plato's articulation of our grasp of the good. Specifically, it highlights the way in which that grasp is implicit in any understanding of the world.

arguments and deploy the tools of classical logic are the lurching first steps along the path of a budding philosopher.[154]

But these strides down the road of study are not really discrete steps that, once covered, one never needs to return to again. Instead, each skill requires iterative practice. Each is then more like a stage in an *activity of discovery*. Taking an elementary logic class, for example, might require months of hard study, and what is valuable about its study, knowledge of its good, is only readily available until long after getting underway. Where at first a logical fallacy in one's thought might pass by unseen, the trained philosopher's eye latches on to such gaps in argument, perhaps even a cultural and historical prejudice that lies at its root, thus seeing in clear and distinct focus what may have been invisible before.

What else can we say about this case? We know that discovering the value of philosophical study is not always easy. In fact, it often takes years of dedicated work before one can even begin to make contributions of one's own. Like any other activity, the philosophical novice must learn the tools of the trade. Doing so requires acquiring the conceptual tools of argumentative analysis, learning to deploy those tools in truth-preserving ways while also developing an increased sensitivity to what is already familiar. Being able to describe what is happening, even in our own minds, requires quite a lot of disciplined, creative work. Minimally, requires *attention*, both to the acquisition of the conceptual tools of the philosopher, but also to the structure of philosophical writing, the art of observation and most importantly, to the world and one's own prejudices. Moving from the philosophical novice to a competent student and then

[154] As we saw in the previous chapter: learning is a practical *achievement*. It is easy to mistake the good in this case. What is difficult in the case of professional philosophical study is to avoid egoistic pursuit of one's own ideas, but rather to stay focused on uncovering the difficult reality of the phenomena in the world Murdoch shows us that the self can get in the way of grasping a difficult and elusive value. She calls this conceptual obstruction a flight into fantasy.

on to a professional in the field requires a developing attentive sensitivity that is an iterative, ongoing process of discovering the good, one that is as much about finding new ways to imagine the familiar as it is about the disciplined adherence to philosophical method.

To continue down the path of a philosopher requires a special *stance* toward philosophical inquiry itself. Where curious intrigue might be enough to begin a provocative introductory text, it is not always enough to sustain the novice through an entire introductory class. Moreover, students new to the field often have a difficult time separating attacks on their *beliefs* from attacks on *them*. When pressed for justification for their views about perception, God or morality, they tend to see it as a threat to personal attack on their intelligence or character. Developing the attitude of open, honest and diligent collective inquiry that is, when done well, aimed only at finding and championing the truth, takes practice and habituation from the norms of debate that may be familiar in other walks of life. Seeing what it takes to unravel our web of beliefs, assumptions and habits, how to make them more consistent and why it is important to do so requires more than just a curious mind, more than attentiveness to the arguments of great philosophers; it also involves developing a *love* of philosophy.[155] The stance of love is necessary to see philosophical study as ultimately a collective enterprise, aimed at truth rather than acceptance, humble questioning rather than a system for dogmatic attack. Ultimately, then, seeing what is worth going for in philosophy is a matter of falling in love with it.

Recall the example of burning cats. Loving attention in this case would open up to a fuller meaning of the scene: what is heinous about this case not merely the harm done to the cat, although that is significant. Part of the horror is also to imagine the psychology of children capable of such abuse; to understand the deformations of character that they are expressing and

[155] Indeed, *love* is right there in the word *philos*.

deepening. A loving attentive gaze is one that highlights not merely the apparent transgression, but also the pathology that surfaces in this case. Loving attention helps envision the whole scene, to see the depth of the evil in all its significance.

As a model of practical thought, loving attention looks even more attractive when we turn our attention to the appreciation of the value of individual human beings. We only learn the value of particular others in and through direct experiences with those individuals. While it is true that we can, without any direct experience, say that we have sufficient reason for affirming the value of human beings *in general* and thus acting from these convictions *here and now*, it is quite another thing to assume that we can actually obtain the evaluative attitudes constitutive of these general commitments through argument or consideration of reasons. As we have seen in chapter one, this is not always the case. For, we do not come to love someone through careful consideration of our reasons for doing so. Instead, the good of an individual unfolds to us from the dialectical process of getting to know *them*, in the fullness of their inexhaustible complexity. No argument could possibly move one to love, nor provide any sort of grasp of the value that we appreciate through loving. To assume that one could, without direct experience, come to understand the value of other human beings through the consideration of reasons, or even justify our love in this way, is to accept the substitutability of our experience and our beloved. Kyla Ebels-Duggan makes this point clearly:

> The idea that all of our commitments, and most especially our normative commitments, should be subjected to critical scrutiny, and the stronger idea that we should maintain only those that we are able to justify on reflection, is often presented as a guiding

> philosophical ideal, even as constitutive of doing philosophy. But I claim there is no stance or attitude that would count from the point of view of, for example, a parent as appropriate, yet genuine, critical consideration of the question whether her child is, in truth, valuable in a way that makes her worth loving. (Ebels-Duggan forthcoming, 16)

The point here is not merely that we do not normally reason our way into love, nor that calls for reasoned justification for such love are difficult to parse, but more importantly that the question of reasoned justification cannot even be taken seriously from within the loving perspective that is being challenged. This is because there is a fundamental methodological schism between what philosophical argument on the basis of reasons can do and the attitude of love.

Suppose your colleague just bought a new car and is telling you how lovely it is to drive, how it beats riding a bicycle. Now suppose you call your colleague's belief in the superiority of cars as a mode of transportation into question. In so doing, one could imagine a line of reasoning that traces the historical origin of automobiles in America, notes their devastating effect on pedestrians, neighborhoods and unpacks the racialized placement of highways and roads to facilitate the movement of white elites into and out of cities. All this could represent cogent reasons for rejecting that initial belief, and one could easily imagine your colleague following you to your conclusion and changing her belief. So much looks like the paradigm case of the power of successful critical thinking.

At first glance, the parallel seems to work for falling out of love. If one were given damning evidence that one's beloved is, say, a compulsive liar, it seems plausible that such reasons for rejecting one's feelings in this case can move us. But here an interesting asymmetry arises when we consider falling *in* love. Since the attitude in question requires appreciation of the

value in question, which can only be had via experience of its object, it follows that no one could *deduce* their way into love. Indeed, to treat the love of a human being as a possible object of intellectual curiosity is to alienate oneself from the intimacy of the position of the loving relation itself, which is the only position from which appreciation of the relevant value is possible. Justification of such an attitude cannot be had from any public point of view because the interlocutor who questions its value cannot even be taken seriously from within that relation, let alone offered a way to see through to that value, absent the relevant experience with it.[156]

Loving attention looks more promising when we move to talk of valuing individual human beings because it does not attempt to explain our attitudes and commitments like love by any appeal to argument or the weighing of reasons.[157] Instead, my view makes sense of love for an individual is that it amounts to first-hand appreciation of the value of their beloved. This requires an attentive gaze that proceeds stepwise, gradually opening wider vistas of possibility for who that person is and what they are capable of. Such widening and deepening of appreciation is surely both a perceptual and cognitive capacity, both passive and active aspects of our imagination that build up, through time and sustained engagement, our knowledge of that person. In this sense, our love is grounded by an awareness of the value of the beloved, rather

[156] As I noted in chapter one, love is a sensible attitude to speak of in general, and its content (what we do in its name) is subject to intelligible challenge. But this is not the same as claiming that the value of the person is equivalent to the generic reasons offered in defense of what we do in their name. That is to substitute the generic relation description for the value itself.

[157] Many contemporary philosophers simply *define* the moral outlook as in direct tension with partial, personal endeavors like love. See, for instance, Railton (1984), Williams (1981a), Wolf (1982), Scheffler (1994), and Frankfurt (1988). On their views, the moral outlook is a disinterested one, free of any subjective commitment to any one individual, including love, which might cloud our ability to reason objectively. Entering into moral deliberation necessarily involves considering only those impartial demands that bind us all. But to make such a division is to accept a version of the formal approach that we rejected in chapter one. By ruling love of individual human beings out of the scope of morality from the start, we are avoiding, not answering the issues about practical reasoning raised in our discussion.

than any inference or argument. Loving attention is a dialectical process of revealing reality, which subsequently brings one's reasons into view.

What this view reveals, if nothing else, is that practical thinking can involve more than the episodic weighing of reasons. Loving attention highlights the way practical thinking can also be a continual activity that begins and develops by attempting to bring the world more clearly into view. As we have seen, this clarification is best described as an activity because it is an ongoing process that must be actively sought out, and rather than terminating in an action, is always ongoing. The importance of this process cannot be stated for all to see, but instead requires time and effort to uncover for oneself.

V

There are a number of objections that need to be addressed. The first involves the way I have characterized the relation between value acquisition and reasons. Consider the philosophical novice. Where I have said that the novice is *ignorant* of the reasons of the expert because the former lacks the necessary experience that partially defines the latter, one might say instead that the two simply apprehend *different* reasons. On this view, we might explain the inarticulateness of the novice in terms of their contingent conceptual limitations. After all, part of becoming initiated into the practice of philosophical inquiry is to practice seeing these initially hidden concepts and considerations that one unaccustomed to the field does not immediately grasp. In this sense, the philosophical novice and the expert might be said to encounter different reasons. Where the novice can speak in vague generalities, making reference to what they find interesting or curious, the expert can speak with more mastery, stating directly the power and sustaining fountain of energy that comes from deploying philosophical arguments to clarify the reality of

the world. Where the novice *hopes* for such experiences, the expert *knows* them. As a function of their differential experiences, their different locations along the path of study, the novice and the expert simply have different reasons for going on.

There is something right about this objection. There are material differences between what it makes sense for the novice and the expert to respectively do. As a result, we should expect that each has different reasons to follow their respective paths of study. Consider a college freshman taking their first philosophy class (novice) and a philosophy professor (expert). For starters, the novice and expert read different texts. Where the expert might read Kant in the original German, the novice has to read introductory texts. Moreover, it is *good* for the novice to read the considered interpretations of experts before, say, venturing alone into the murky waters of the *Critique of Pure Reason*. Similarly, it would make little sense for the expert to settle for the interpretations of others when attempting to reconstruct their own. In this case, the novice and the expert *do* seem to be responding to different features of philosophical inquiry corresponding to their different levels of sensitivity to the details of argumentative analysis. It seems, then, that perhaps the two really are responding to different reasons: where the novice has a reason to pursue introductory texts for the sake of simplicity and clarity, the expert has reason to avoid these texts.

Despite its initial attractiveness, I think that this way of parsing the appreciation of activities like philosophical study lacks the essential unity of our phenomenal experience. According to the objection, what is in fact an ongoing, continual process of coming to appreciate an object of value appears to be a series of disconnected episodes, each disjointed and isolated from the rest. This view feels at home within a picture of practical thought whereby the

paradigmatic conclusion of a practical deliberation is the adoption of an intention to act.[158] But as we know, this is not necessarily what is going on when we come to appreciate the intrinsic value of philosophical study. Surely we *do* adopt intentions to act along the way at any given moment (e.g.,, going to class, picking up a book, reading it carefully). But if we equate the intrinsic value of the activity of philosophical study with the series of individual intentional actions necessary to do so, we mistake the *mere motions* for the thing itself. In fact, I think such a view passes over the relevant value in question entirely. What is valuable about philosophy is *not merely* the sum of the series of one-off intentional actions that lead to reading or writing its constituent literature. While we *must* form intentions in specific episodes of practical deliberation in order to appreciate philosophical study, such appreciation is not *reducible* to the sum of the conclusions of episodic deliberation that yield such actions; the latter are merely the disjointed generic descriptions that approximates the experience of the philosophical student from the outside looking in.

Consider a mathematical analogy. A Riemann sum is an approximation of an integral by a finite sum. A Riemann sum measures the area under a curve by partitioning the curve into a finite series of shapes (usually rectangles) and adding all these smaller areas together. The smaller and closer together the partitioning shapes, the more accurate the area estimate. However, since a curve is not equivalent to a series of rectangles, A Riemann sum will always be an estimate. And while it can be accurate enough for certain purposes, to suppose that it *just is* an integral is to confuse two entirely different mathematical concepts, to mistake an approximation by the sum of individual geometric partitions for the unified curve itself. Just as a Riemann sum is not equivalent to an integral—but rather an approximation based on the collection of distinct points

[158] This dichotomy between two pictures of practical thought is from Brewer (2011). The picture I only briefly sketch here is what Brewer calls the modern understanding of practical thought.

along the way—the intrinsic value of philosophical study, or the value of a particular human being, is not constituted by a finite series of episodes that culminate in intentions to act on the basis of reasons. While learning to appreciate the intrinsic value of philosophical study surely *can* yield new reasons to be offered to others, these are two distinct, albeit related, things. The appreciation of philosophical study, like the integral curve, is more than the sum of a series of finite events that mark its approximation.

In response, another objection arises: why not think that the approximation of practical thought in terms of reasons *is good enough for morality*? Do we really *need* more than a close approximation of the value of philosophical study or the love of a particular human being to get on with the business of building moral theories? One might agree that the value of philosophy or human beings cannot be exhausted by talk of reasons available to the rational intellect independent of experience. However, in the case of interpersonal moral justification it is trivially true that we can never have perfect knowledge of each other. That is partly why we *need* moral theories. So, on this view, it would be unreasonable to demand that we need full knowledge of the intrinsic value of every particular activity, experience or person in order to construct moral principles and apply them to specific cases. In other words, the value we need to see in all human beings in order to grasp the content or the seriousness of our moral obligations is not the particular, idiosyncratic sort of value that I have focused on.

Many philosophers see things this way.[159] Following Scheffler, you might think that a moral theory is nothing but "a system of general principles of moral evaluation that, when combined with the relevant bits of specific information, generates accurate overall moral verdicts

[159] Hare (1963; 1981) is perhaps the fiercest defender of this way of looking at what moral theories are.

about particular actions."[160] This is a version of the substantive approach outlined in chapter one. On this approach, we can rule out the importance of appreciating the full value of philosophy, art, nature, or human beings for doing moral theory by specifying some substantive conclusions about the scope and point of moral theory. On this view, all we need is knowledge of what is relevant to the domain of moral life, together with a set of principles for sorting and weighing such relevant features and we are off.

I think that there are at least two reasons to reject this approach. The first is that my account of loving attention does not require that we have perfect knowledge of every person's experience in order to make fair moral judgments. Instead, my claim is that this way of looking at what morality is and cannot fairly represent *some* familiar and important activities in life. For example, practical experience with instances of loving relationships puts us in touch with the sort of irreplaceability that is basic to morality. If such values include, as I have argued, the appreciation of individual human beings, the very value claimed by many to be at the *heart* of moral philosophy, then it appears that moral theorizing has become a seriously protracted enterprise. This exclusion might seem appealing. Indeed, one might insist upon the idea that there are certain moral demands (perhaps absolute prohibitions) that hold for everyone for all time. Moreover, one might hold that we can know what they are without a potentially endless review of the idiosyncratic considerations that might, if taken to heart, induce people to think they are not bound by these imperatives.

Despite its appeal, I think there are a number of difficulties with this approach. First, it fails to explain how we could justify such pursuits to others when those pursuits conflict with other demands, nor does it provide a way to fairly assess such pursuits against other important

[160] Scheffler (1992), 39.

moral demands. This is a serious charge, since assessing the priority of moral demands is part of the central work of any moral philosophy. Second, it does not explain the phenomena of some very basic aspects of life, such as why and how we come to value that which we do not yet have the words to describe or understand, nor does it explain how such activities could have any resonant unity through time. Simply denying the importance of specifying the value of these pursuits in reasons is to ignore, not defend or explain their exclusion from what is significant in or relevant to moral deliberation.

But even if we grant that loving attention *does* name a distinct form of practical thinking and grant that its cases *are* morally relevant, the objector might still hold that my view does not solve this problem about interpersonal justification. Recall that one of the central concerns of my account, indeed an identifying characteristic of this mode of practical thinking, is that the appreciation of certain kinds of values simply *cannot be fully appreciated by those outside its practice*. Suppose one were to grant this point: it is unclear what moral theorists should do with this information. If, on the one hand, theorists agree that some values cannot be properly weighted in interpersonal justification—cannot be accurately captured in terms of reasons—then the entire enterprise of moral theory looks problematically incomplete. If, on the other hand, theorists attempt to accommodate such values, to shoehorn them into the picture of practical thinking that takes reasons to be the entirety of interpersonal justification, then it seems that theorists are doomed to misrepresent these values in.

Doomed? Recall our argument from chapters one and two. I claimed that there are sources of resistance to description in terms of reasons, whose value cannot be fairly assessed in the modes of assessment that moral theory tends to insist upon (reasons for action). As a result, the prospects for elaborating any set of principles that meet something in the neighborhood of

reasonable rejectability appears bleak. We saw that this is because some powerful objections will have to be dismissed out of hand on the ground that they emanate from a practical appreciation that has no place in the assessment of general moral principles. To be sure, this is no *decisive* objection to theories like Scanlonian contractualism and its nearby allies, it does bring out what seems to be an unavoidable weakness in theories of this kind. Specifically, it unveils an unavoidable table-thumping on behalf of the conception of practical thought in whose terms they are elaborated.

I think a full resolution to this problem is beyond the scope of the current discussion, but I will offer two brief remarks. First, recall that the present discussion began with a claim about the inadequacy of the dominant methodology for building moral theories. If the result of my project is to vindicate just this claim, I will be satisfied. If one found that initial claim well founded, or at least plausible, then one need not *demand* that a fuller description of this inadequacy yield a solution that fits neatly within the bounds of mainstream moral theorizing. What to do about this larger problem is a separate question. This chapter is only a first step toward an alternative; it is by no means a complete revisioning of the methodology of moral theory building in general.

We can also say *something* about how moral philosophy ought to proceed. One might think that the account of loving attention offered here does not contradict, but rather sits nicely with a version of particularism. As we saw in chapter two, the basic tenets of particularism (holism and anti-codifiability) dovetail with the idea of categorial generality. If we think that appreciation of the intrinsic value of certain category of activities can only be had through direct practical experience with its instances, then it seems plausible that such values must also be analyzed holistically, since the circumstantial details that will determine the moral significance of

the case are *to be discovered*, not stipulated in advance. For this same reason, one might argue that the results of moral deliberation about these kinds of activities are not codifiable. If the appreciation of a category of value requires experience with its instances, which are unpredictably variable and complex, then it seems plausible that we will not be able to codify the significance of that value outside its experiential context. If we could, then sufficient appreciation of certain values actually would not require first-hand experience. While surely just a sketch at this stage, this connection has the potential to develop into a compelling argument for particularism.

Consider one final objection. Suppose one concedes all three points above: that my account of loving attention (i) better describes the phenomena of the cases in question than the dominant methodology of representing that value in terms of reasons, that it (ii) accounts for values that are necessary for building plausible moral theories, and (iii) provides a compelling case that the methodological assumptions of modern moral theory building need revision. These are large concessions. But even so, it might still be argued that what I have presented here *is not really an account of practical thought at all*.

Instead, it looks as though loving attention is a *hybrid account*, which conflates instances of practical thinking with episodes of accompanying theoretical thought. Where the traditional literature marks the boundaries between theoretical and practical thought precisely at the temporal start of deliberation and the adoption of an intention to act (or perhaps the resulting action) I have drawn the boundaries of practical thought much differently. On my view, it is difficult to tell *what is and is not practical thinking*; my view makes it seem as if practical thinking goes on at every moment, lumping together seemingly obvious cases of mere *theoretical*

thought, like the acquisition of moral concepts through reading, alongside clear cases of *practical* thought like forming an intention to pick up a book. As a result, so the objection goes, my account represents neither an objection nor an alternative to the traditional account practical thinking, but rather a conflation of fundamentally different types of thought.

It is true that my account does not mark a sharp distinction between theory and practice. However, I believe, perhaps unsurprisingly by now, that this blurring of the boundary between the two categories is actually an advantage of my view. We have good reason to reconsider the traditional distinction between theoretical and practical thought once we consider a point that has been developed earlier. Careful consideration of the acquisition of moral concepts suggests that the process through which we gain the language to describe and act on the world is much more complex than is usually discussed. Specifically, such concept acquisition can itself involve extensive practical experience within a particular social and historical context.

For instance, while it is possible to learn the definition and basic criteria for deploying the word *honor* through, say, voraciously reading war novels, it is only through inculcation into a set of historically contingent cultural practices that one can comprehend the sense of *honor* that takes being prepared to die for one's country to be a high ideal rather than a grim tragedy. Such an understanding of the term stems from a collection of ways of life that elevates the desire for honor beyond a mere superlative (e.g., as in an *honor roll*) to a lifetime achievement (e.g., "he was the most *honorable* man I knew"). As such, deliberation on what is *honorable* to do necessarily involves quite a lot of experience and thought, which together form the foundation of the live options from which any practical choice is made. The evidence for this is the wide variability in the meaning of the term across culture and time. For some honor is the highest glory in life, where for others it is merely a synonym for gratitude. At every moment of

deliberation, we are always already drawing up from our well of background experiences, deploying concepts formed from previous episodes that form our sense of what life is like. It is only in and through an extended collection of practical experiences (living with others, joining youth organizations, venerating *these* people as heroes and not *those*) that together form the cultural well from which a robust understanding, that one can draw on that more embedded sense of *honor*. As a result, I think a blurry boundary between theoretical and practical thinking, while inexact and perhaps less theoretically satisfying, is both more reflective of the phenomenal quality of lived experience and better positioned to spark a richer discussion of moral deliberation and the appreciation of value than the current philosophical literature on reasons permits.

VI

This chapter sketched three advantages of my account of loving attention as a mode of practical thinking: that it (i) better explains the phenomenology of a certain familiar mode of valuing, (ii) suggests an underexplored argument in favor of particularism and (iii) provides a messier, but ultimately more accurate understanding of the category of practical thinking. I want to conclude by discussing one more consequence of my view.

Where talk of actions done for reasons prompts philosophers to reach for the metaphor of *weighing*, my account of loving attention lends itself to the metaphor of *depth*. Weight, as a measure of things, is an abstraction; it is an attempt to bring disparate objects under the same scale for comparison. For instance, apples and cell phones do not share much in common. However, by abstracting away from their observable features, we can use their weight as a common scale upon which to compare. Reasons prompt such thinking because of their function

as the central currency of moral justification. The thought is that deliberation often requires us to compare the worth of actions as disparate as apples and cell phones. *Whatever the disparate values in question*, reasons can be the common scale for comparison. The *weighing* of reasons is thus a favorite metaphor of philosophers, in part, because it frames moral deliberation as an objective enterprise, like Lady Justice feeling the scale of reasoned argument tip.

We now have some evidence that this way of thinking about moral deliberation has a few serious drawbacks. For one, chapter one argued that not all considerations are so easily abstracted from such that their proper *weight* can be assessed from the public point of view. But beyond the possibility that some reasons may not be fully appreciated (weighed improperly) from the point of view assumed by this way of thinking about moral deliberation, we have seen that some values are difficult to *get on the scale at all*. To stretch the metaphor, Lady Justice may be *blind to them*. Significance aside, how do I even *speak* of my love for my partner in the same register as, say my desire to file my taxes? While I might clearly see obligations flowing from both, and indeed even be able to vaguely point to some set of concrete tasks demanded by each, it is difficult to imagine how anyone could take seriously the suggestion that these two sources of demanding value can be meaningfully *registered* and *weighed* against each other.

In contrast, my account suggests a more insightful picture of what is going on here. Rather than attempting to flatten all human concerns into a metaphorical abstraction like weight, moving beyond talk of reasons allows us tosee that my beloved and my taxes are sources of value that make practical demands at different *depths*. Where taxes do clearly elicit my practical thought, the sense in which they do so is only *skin deep*. In contrast, when my partner comes home from a long day at work, my sense that she needs to vent takes clear precedence. The question is: why? While one certainly *could* describe my choice to respond to my partner's needs

over finishing my taxes in terms of one object of attention providing a *weightier set of reasons* than the other, to do so would be to suggest that these two objects exist at the same level, to use my metaphor, at the same *depth of life*. Reasons talk makes my choice to focus my attention on one rather than the other look as if it were predicated merely on consideration of the abstract *weight* of reasons on each side.[161]

We then considered what I called the *standard view*, which relies heavily on the articulation of generic reasons for action to model what is worth going for in a given situation. But at this stage we can note that this example is not easily captured in terms of reasons for action; we do better if we draw this example as a call for loving attention that is rooted in the rich, particular history of our lives together. Such a description forms the backdrop and cultural well from which I draw my energy to instinctively put down my W2 and make my partner a cup of tea.[162] As taxes and partners go, so follow their corresponding modes of practical thought (actions done for reasons and loving attention): where one runs *shallow*, the other *runs deep*.

Let us conclude by comparing my analysis of conflicting demands to Scanlon's. This will allow us to neatly recapitulate what has been gained by the present discussion. Scanlon asks us to imagine a game of tennis where one must decide whether to play to win or go easy on the opponent so as not to inflict any hard feelings. After deciding to play to win, Scanlon writes,

---

[161] This is a similar, but slightly different point than Williams' (1981b) "one thought too many" case. Whereas for Williams' *any* deliberation upon whether or not to save one's drowning wife represents a moral failing, my point is that the explanation for why one need not deliberate is that love and, say, the demands of strangers exist at different levels of depth in life. Such demands rarely cross paths in the way imagined by reasons talk.

[162] Lawrence (2018) notes how Philippa Foot speaks eloquently about the connection between *depth* and *happiness* in human life. Lawrence notes that for Foot, "happiness in this deep sense is one of a human adult, and in its nature impossible for a child, it is not detachable from an adult's "resources of experience and belief" but something that "must extend all the way into the under- lying thoughts that a person has about himself and his life." For Lawrence, "The proper objects of deep happiness are [what Foot 2001 calls] the "basic" projects of life.. So important is this task, that it marks the "expression and exploration of our humanity." (Lawrence (2018) 201).

> [r]eaching this conclusion involves deciding which reasons will be relevant to how I play. The fact that a certain shot represents the best strategy will count as sufficient reason for making it. I need not weigh against this the possibility that if the shot succeeds then my opponent will feel crushed and disappointed. This does not mean that I cease to care about my opponent's feelings. I may still want him to be happy, hope that he is able to take pleasure in the game, and refrain from laughing at his missed shots. My concern for his feelings is not eliminated or even diminished; I just judge them not to be relevant to certain decisions. (Scanlon 1998, 52)

Scanlon's point is that there are cases where conflicting reasons need not be weighed because we have a further, higher-level reason *reorder* the reason-giving force of certain considerations, to deem certain one more or less relevant in these circumstances. Just as tennis players can bracket their opponents' feelings in the heat of the game, teachers must sometimes have reason to bracket considerations about their own interests for the good of their students. He concludes that the reasons we have for living up to the standards associated with such roles are reasons for reordering the reason-giving force of other considerations: reasons for bracketing some of our own concerns and giving the interests of certain people or institutions a special place.

Scanlon is surely right that we often need not appeal to any weighing procedure in order to resolve conflicting considerations. But I think his analysis gets the phenomenon wrong. Consider my frustrated wife once more. On Scanlon's view, upon realizing that I have conflicting considerations (finishing my taxes and caring for my wife) a new, higher level ordering consideration appears to reorganize my priorities and point the way to the tea kettle.

This cannot be right. In this case, it is not that my desire to finish my taxes triggers an appeal to any higher-order consideration at all. Instead, my suggestion was that such a desire is best described in terms of depth: whatever reasons I have to finish my taxes could never conflict with the reasons I have to care for my wife because they operate in different registers.

What does this talk of *registers* mean? By breaking the bonds of reasons talk, we gain insight into the different ways we value the diversity of goods in life. Where Scanlon's insistence on reasons-fundamentalism places the motivational force in cases like this in an alien, overly intellectual frame, my view need not generate any higher-level reason in order to explain the resolution of conflicting considerations. Where talk of reasons flattens the terrain, collapsing all forms of valuing into reasons-giving considerations, we would do better to embrace the diversity of ways we can apprehend good in the world, ways that operate at different levels of attachment and meaning, which may or may not be capable of being compared or weighed on a common scale. In moving beyond talk of reasons, we adjust our eyes to better apprehend the profundity of what is valuable in life.

Conclusion

I argued that mainstream philosophy fails to account for the importance of some phenomena that are clearly central to moral life. The claim was that why we pursue energeia such as philosophical inquiry, art and loving relationships with particular people cannot be heard in the language of reasons without distortion. This is because that language holds a high standard of articulacy as a barrier to any consideration being granted the status of a genuine reason. We saw that such a standard cannot be met in cases where the value of the activity, experience or relationship in question cannot be clearly understood, much less articulated (by the agent attempting to encounter that value or by anyone else) except through a period of wholehearted engagement with that very pursuit. This means that what is worth going for in these kinds of pursuits resists articulation in the form required by moral theory: an individuated, generally stated reason. But if these are the only grounds for practical moral thought, then the result is an artificial restriction on the kinds of considerations that do and do not count, one that leaves out a proper understanding of these central and underexplored class of phenomena.

Chapter one was an investigation of two different methods of theoretically restricting the domain of genuine reasons. Its aim was to highlight what I see as a serious oversight at the heart of modern moral theorizing. In particular, there exists a class of familiar and important human values that resist being brought under the language of public deliberative exchange that characterizes theoretical attempts to model reasons in moral deliberation. I called them energeia. What falls under this description are not esoteric whims, but instead deeply familiar human activities: appreciating art, philosophy or particular human beings are just some examples of activities whose appreciation is irretrievably tied to first-personal experience. Since these are allegedly *central* values in the study of ethics, it is cause for concern that their full appreciation is

either ignored or downplayed by the methodology we have considered. Since the reasons that justify these activities are only available for true appreciation after undergoing a first-hand attempt to undertake them, they are not fairly represented (if shown at all) in philosophical models of moral deliberation that take publicity or universality to be a fundamental attribute of reasons.

Chapter two turned to a prominent version of particularism as a particularly appealing alternative approach to capturing what is valuable about these pursuits. Specifically, the denial of general rules for specifying right action, combined with holism about reasons and the anti-codification thesis seemed to offer a much more promising solution to this restriction problem. Rather than searching for some general principle about when a consideration counts as a reason, or when a consideration is morally relevant, particularism suggests that whether studying philosophy is a source of genuine value will be a function of a holistic evaluation of the situation in which the agent encounters the discipline. This is a much more rich and compelling approach because it better fits the individual, historically contingent experience of coming to be initiated into a discipline, activity or relationship. Whether and to what extent philosophical study is valuable does, indeed, seem to be a function of a whole host of factors that escape general codification in advance of a period of attempting to bring its value into view.

However, we saw that particularism's most prominent view, Jonathan Dancy's, fails to explain our central phenomena because it is grounded in a picture of practical thinking that contains the same philosophical prejudices we saw in chapter one. Dancy's account fell short because of its commitment to what I argued was an over-intellectual idea of how we encounter reasons in moral life. Specifically, we saw that the claim that practical thinking is a two-stage affair, where what grounds any particular action as justified is the prior encounter with some

given, individuated consideration of that particular situation. This general gloss seemed in tension with holism about reasons, the anti-codification thesis and most importantly the phenomenology of moral life we considered in chapter one. What was missing, then, was a vision of practical thought that could make sense of the phenomena that began our investigation, one that harnesses all the power of a particularist approach without the philosophical baggage of a Dancy-style account. Filling this void was the task of chapter two.

Chapter three took up the work of early philosophical work of Iris Murdoch with the aim of sharpening and clarifying account of the phenomenology of moral deliberation and concept acquisition. Narrowing in on her discussions of activity, deliberation, imagination and loving attention, we saw that while Murdoch's view is certainly sympathetic to the general outline of the particularist project, her work has a much wider application. In particular, we saw that her account is best understood as grounding a much wider, more complex account of agency in practical thought and moral deliberation.

Building off the work of Iris Murdoch, the final chapter sketched a positive account of practical thinking that I called loving attention. My claim was that capturing why we pursue energeia is best done not in the language of reasons but rather in the language of love. With an account of Murdoch's philosophical work on love, attention, deliberation and the experience of moral life on the table, we then deployed some of Murdoch's conceptual tools to sketch the mechanics of my view, showing how thought makes itself practical in the kinds of cases that have been our central concern. This required a reexamination of those cases with fresh eyes, comparing the explanatory power of my view to the ones we have already discussed. I then defended my account from important practical and theoretical objections. I argued that while

loving attention may be less theoretically tidy, it does a better job of capturing the phenomena of appreciating values in cases where reasons talk falls short.

I hope to have shed some light in a handful of philosophical dark alleys. First, to put words to a deeply held, but often overlooked commitment to genericity in the methodology of moral philosophy. This commitment takes many forms, but the focus of our discussion was in the way the dominant conceptions of what a normative reason must be works to exclude certain forms of valuing from fair consideration in individual moral deliberation and interpersonal justification. If reasons are the currency and values are the goods, then there are goods in the marketplace of moral life that cannot, in principle, be assigned a value in that currency. And if reasons are taken to be the entirety of what makes an action right or wrong, but they do not capture the whole of what it is we value in some important and central cases in life, then there is a serious and unacknowledged limitation on the completeness of moral philosophy.

Second, in uncovering this shortcoming of philosophy's commitment to generic representation of particular values—and its subsequent limitations embedded in its background assumptions about practical thought—this discussion offered a new and potentially powerful argument against the use of general principles. Through consideration of Dancy's work, we saw that the root of the issue of unjustly excluding an underexplored class of considerations (that I called energeia), the problem identified in chapter one, was not the content of any particular generalist theory, nor the *generalizing impulse itself*. The heart of the issue is a widespread set of methodological commitments about how reasons connect to action, how we encounter them, and the conception of agency that combine to artificially limit the scope of practical thinking are left intact by the particularist.

This limitation also represents a missed opportunity. After all, we have seen how holism about reasons and anti-codifiability seem to produce a clearer picture of how practical thinking works in the cases we identified. Where the clear articulation of reasons does play a central role in most instances of moral deliberation, we saw in chapter one that reliance on reasons talk cannot account for cases where appreciating the value of a person, experience or activity is inseparable from the first-personal attempt to bring its value into view. Holism about reasons offers an explanation of why reasoned articulation of what matters in these cases is at best a meager approximation: because full understanding of the moral character of the consideration in question can only be had through a holistic evaluation of the context in which it occurs. Moreover, allowing that the ways in which a consideration functions within that holistic evaluation cannot be codified in tractable rules explains why the articulation of the value one initially saw in such activities rings hollow. This leads us to a surprising conclusion: the phenomena we highlighted from chapter one—studying philosophy, appreciating art and loving individual people—seem easier to explain on the particularist view. As such, their exploration seems ripe to serve as compelling objections to generalism. Yet what we saw in this chapter is that Dancy (particularism's most prominent advocate) not only misses this opportunity, but cannot himself account for these centrally important phenomena. Third, my account of practical thinking, which I labeled *loving attention*, makes better sense of the phenomenon of individual valuing in cases where its representation in terms of reasons falls short.

How so? In cases where articulable considerations that capture the full force and form of a particular value are unavailable or impossible absent some direct, first-personal practical experience, we do better by speaking of the ongoing, dialectical activity of loving attention. Such activity begins with the questioning of the brute-ordering concepts from which reasons (as

in-virtue-of concepts) are called down. This deeper level deliberation is itself a process of imaginative discovery with passive and active components. Imagination involves a developing sensitivity, a coming alive to bits of reality previously hidden. It also involves a creative activity of ever-expanding categories and possibilities for human life, a gradual building up, through practical experience, of new moral concepts. This imaginative process requires practical experience because both its aspects do—that is, both the expansion and development of a reliable sensitivity to new details of the world and the active creation of new categories and possibilities under which to understand and deploy them—require lived practical experience. In this sense, we saw that *doing runs ahead of our concepts*.

This imaginative process requires a multifaceted dispositional stance I called *love*, whose facets included an (i) an open attentiveness to the object, (ii) a development of a perceptual capacity to see the object of appreciation for what it is and what it could become, and a (iii) cognitive capacity to recognize and deepen one's understanding of the object of attention and therefore the proper treatment of it. Call this iterative process of discovering value, complete with brute concept acquisition and dispositional development, the imaginative work of *loving attention*. Loving attention explains the sense in which the work of practical thinking can sometimes be not to individuate some pre-existing reasons that exist independently in the world and hold generally for all, but rather to bring one's reasons into reality.

What this view reveals, if nothing else, is that practical thinking can involve more than the episodic weighing of reasons. Loving attention highlights the way practical thinking can also be a continual activity that begins and develops by attempting to bring the world more clearly into view. As we have seen, this clarification is best described as an activity because it is an ongoing process that must be actively sought out, and rather than terminating in an action, is

always ongoing. The importance of this process cannot be stated for all to see, but instead requires time and effort to uncover for oneself.

Finally, loving attention as practical thinking provides new insight into the structure of thought itself. It suggests that the standard divide between theoretical and practical thinking is less useful than it appears. Aristotle, for instance, distinguishes them by their different aims. For him, the work of theoretical thinking (*theorētikē)* is unchanging truth, where the work of practical thinking is right action (*prohairesis*).[163] But if we take seriously the suggestion that practical thinking aims at the recognition of a truth, namely, the reality of the world, and theoretical thinking actually prepares the conceptual ground for the possibility of action, the two begin to bleed into each other. I have outlined a way to understand this connection that blurs the sharp boundary set up by the Aristotelian tradition: in deliberation, we are always already drawing up from our well of background experiences. Attempting to understand the world around us involves deploying concepts formed from previous episodes that form our sense of what life is like. Such previous episodes are an extended collection of practical experiences (living with others, joining youth organizations, venerating these people as heroes and not those) that together form the cultural well from which a robust understanding, that one can draw on that more embedded sense of any particular concept. In this way, action prepares the ground for concept acquisition, and concept acquisition opens the space of possibilities for which actions or live options. The two are mutually supportive.

It may be said that the preceding discussion has been part philosophy, part cultural criticism. This is because my philosophical work demanded methodological *groundwork*. My

[163] *NE* 1139a17-b13. Aristotle makes a similar claim about delineating excellences by their characteristic work in *EE* 1218b37-1219a23.

starting observation involved a discussion of a tendency for the dominant concerns of one particular brand of philosophy to deflect us away from the reality of the world. Through examples, I constructed a picture of how the language of reasons for action represents a point of deflection away from the difficulties of apprehending the value of energeia. Such a problem may seem unique to the modern analytic, Anglo-American philosophical approach. This groundwork (uncovering the limitations of our methodological commitment to genericity and reasons) is very difficult from within this approach because we philosophers live and breathe reasons-talk so thoroughly that we sometimes fail to see the restrictions it imposes. To the naked eye, reasons look like philosophical oxygen: invisible, essential but therefore also limiting. Such claims *do* look a lot like cultural criticism.

Let me close by saying that I do not take that label to be a bad thing. If this dissertation has been fruitful, it must be, in part, because it reopens an old argument about what ethics is or ought to be. For my part, there is no more worthy philosophical contribution. And since every philosopher is rooted in a complex web of overlapping and intermixing cultures—complete with a host of historically salient insights and prejudices—normative philosophical argument is always, already a kind of cultural critique. To deny or ignore this fact, when writing a dissertation in the field of ethics, would be to fail to honestly face the complexities of historical reality; it would be to ignore the details of how ideas form, distort and shape real people. With this project, I hope to have succeeded in making a small philosophical contribution. But I hope even more fervently to have done justice to the gravity of its cultural subject.

References

Alvarez, Maria. 2010. *Kinds of Reasons: An Essay in the Philosophy of Action*. Oxford: Oxford University Press.

Anderson, Luvell. 2017. "Hermeneutical Impasses." *Philosophical Topics* 45 (2): 1-20.

Anscombe, G.E.M. 1957. *Intention*. Oxford: Oxford University Press.

Apel, Linde. 2011. "Voices from the Rubble Society: Operation Gomorrah and its Aftermath." *Journal of Social History* 44 (4): 1019-1032.

Aristotle. 2011. *Eudemian Ethics*. Translated by Anthony Kennedy. Oxford: Oxford University Press.

———. 2002. *Nicomachean Ethics*. Translated by Sarah Brodie and Christopher Rowe. Oxford: Oxford University Press.

Benjamin, Walter. 2007. "Art in the Age of Mechanical Reproduction" in *Illuminations*. Edited by Hannah Arendt. New York: Shocken Books.

Berlin, Isaiah. 1978. *Concepts and Categories: Philosophical Essays*. New York: Penguin Books.

Blum, Lawrence. 2012. "Visual Metaphors in Murdoch's Moral Philosophy." In *Iris Murdoch, Philosopher: A Collection of Essays*. Edited by Justin Broakes. Oxford: Oxford University Press, 307-324.

———. 1994. *Moral Perception and Particularity*. Cambridge: Cambridge University Press.

Bremner, Sabina Vaccarino. 2022. "Practical judgment as reflective judgment: On moral salience and Kantian particularist universalism." *European Journal of Philosophy*. 30 (3): 1-22.

Brewer, Talbot. 2022. "An Economy of Life." *Raritan* 41 (4): 1-18.

———. 2011. "Two Pictures of Practical Thinking." In *Perfecting Virtue: New Essays on Kantian Ethics and Virtue Ethics*. Cambridge: Cambridge University Press, 116-146.

Callard, Agnus. 2018. *Aspiration: The Agency of Becoming*. New York: Oxford University Press.

Carson, Ann. 2014. *Eros the Bittersweet*. Princeton: Princeton University Press.

Cleveland, Timothy. 2022. *Beyond Words: Philosophy, Fiction and the Unsayable*. Lexington: Lexington Books.

Crisp, Roger. 2012. "Iris Murdoch on Nobility and Moral Virtue." In *Iris Murdoch, Philosopher: A Collection of Essays*. Edited by Justin Broakes. Oxford: Oxford University Press, 275-292.

———. 2000. "Particularizing Particularism." In *Moral Particularism*. Edited by Brad Hooker and Margaret Little. Oxford: Oxford University Press, 23-47.

Cullity, Garrett & Holton, Richard. 2002. "Particularism and Moral Theory." *Aristotelian Society Supplementary Volume*. 76:169-209.

Dancy, Jonathan. 2018. *Practical Shape*. Oxford: Oxford University Press.

———. 2004. *Ethics Without Principles*. Oxford: Oxford University Press.

———.1993. *Moral Reasons*. Oxford: Blackwell Publishers.

———. 1985. “The role of imaginary cases in ethics.” In *Pacific Philosophical Quarterly* 66 (1-2): 141-153.

Denham, A.E. 2012. "Psychopathy, Empathy, and Moral Motivation." In *Iris Murdoch, Philosopher: A Collection of Essays*. Edited by Justin Broakes. Oxford: Oxford University Press, 325-352.

———. 2001. “Envisioning the Good: Iris Murdoch’s Moral Psychology.” *Modern Fiction Studies* 47 (3): 602-629.

Diamond, Cora. 2003. "The Difficulty of Reality and the Difficulty of Philosophy." *Partial Answers: Journal of Literature and the History of Ideas* 1 (2): 1-26.

———. 1996. "'We Are Perpetually Moralists': Iris Murdoch, Fact, and Value." in *Iris Murdoch and the Search for Human Goodness*. Edited by Maria Antonaccio and William Schweiker. Chicago, IL: University of Chicago Press, 79–109.

———. 1988. "Losing Your Concepts." *Ethics* 98 No.2, 255-277.

Driver, Julia. 2012. "For Every Foot Its Own Shoe: Method and Theory in the Philosophy of Iris Murdoch." In *Iris Murdoch, Philosopher: A Collection of Essays*. Edited by Justin Broakes. Oxford: Oxford University Press, 293-306.

Ebels-Duggan, Kyla. Forthcoming. "A Question of One's Own." In *Normativity and Agency*. Edited by Tamar Shapiro, Kyla Ebels-Duggan, and Sharon Street. New York: Oxford University Press.

———. 2019. "Beyond Words: Inarticulable Reasons and Reasonable Commitments." *Philosophy and Phenomenological Research* 98 (3): 623-641.

Fernandez, Patricio, A. 2018. "The Good Will Be First." In *Oxford Studies in Normative Ethics Volume 8*. Edited by Mark C. Timmons. Oxford: Oxford University Press.

Foot, Philippa. 2001. *Natural Goodness*. Oxford: Oxford University Press.

Ford, Anton. 2011. "Action and Generality" in *Essays on Anscombe's Intention*. Edited by Anton Ford, Jennifer Hornsby, and Frederick Stoutland. Cambridge: Harvard University Press.

Frankfurt, Harry. 1988. *The Importance of What We Care About: Philosophical Essays*. Cambridge: Cambridge University Press.

Gaita, Raimond. 2004. *Good and Evil: An Absolute Conception*. New York: Routledge.

Gauthier, David. 2022. Rational Deliberation. Oxford: Oxford University Press.

———. 1986. Morals By Agreement. Oxford: Clarendon Press.

Goldberg, Samuel. 1993. *Agents and Lives: Moral Thinking in Literature*. Cambridge: Cambridge University Press.

Hansen, Miriam Bratu. 2008. "Benjamin's Aura" *Critical Inquiry* 34 (2): 336-375.

Hare, R.M. 1963. *Freedom and Reason*. Oxford: Oxford University Press.

———. 1964. *The Language of Morals*. Oxford: Oxford University Press.

———. 1981. *Moral Thinking: Its Levels, Methods and Point*. Oxford: Oxford University Press.

Harman, Gilbert. 1977. *The Nature of Morality*. New York: Oxford University Press.

Held, Virginia. 1987. "Non-Contractual Society: A Feminist View." In *Science, Morality, and Feminist Theory*. Edited by M.P. Hanen and K. Nielsen. Calgary: University of Calgary Press.

Herman, Barbara. 2021. *The Moral Habitat*. Oxford: Oxford University Press.

———. 2007. *Moral Literacy.* Cambridge: Harvard University Press.

———. 1993. *The Practice of Moral Judgment*. Cambridge: Harvard University Press.

Hieronymi, Pamela. 2011. “Of Metaethics and Motivation: The Appeal of Contractualism,” *Reasons and Recognition: Essays on the Philosophy of T. M. Scanlon*. Edited by R. Jay Wallace, Rahul Kumar, and Samuel Freeman. Oxford: Oxford University Press, 101–128.

Hopwood, Mark. 2022. "How to Read Metaphysics as a Guide to Morals." In *The Murdochian Mind*. Edited by Silvia Caprioglio Panizza and Mark Hopwood. London: Routledge.

———. 2018. ‘The Extremely Difficult Realization That Something Other Than Oneself Is Real’: Iris Murdoch on Love and Moral Agency. *European Journal of Philosophy* 26 (1): 477-501.

———. 2017. "Murdoch, Moral Concepts, and the Universalizability of Moral Reasons." *Philosophical Papers* 46 (2): 245-271.

Hutcheson, Francis, 1971. *Illustrations on the Moral Sense*. Edited by Bernarad Peach. Cambridge, MA: Harvard University Press.

Jackson, Frank, Pettit, Philip & Smith, Michael. 2000. "Ethical Particularism and Patterns." In *Moral Particularism*. Edited by Brad Hooker & Margaret Olivia Little. Oxford University Press. pp. 79--99.

Korsgaard, Christine. 2009. *Self-Constitution*. Oxford: Oxford University Press.

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

Larmore, Charles. 2021. *Morality and Metaphysics*. Cambridge: Cambridge University Press.

Lawrence, Gavin. 2018. "The Deep and the Shallow." In *Philippa Foot on Goodness and Virtue*. Edited by Micah Lott. London: Palgrave-Macmillan. 187-255.

Levibond, Sabina. 2020. "The Elusiveness of the Ethical: From Murdoch to Diamond." *Royal Institute of Philosophy Supplement*. 87: 181-200.

Little, Margaret Olivia. 2000. "Moral Generalities Revisited." In *Moral Particularism*. Edited by Brad Hooker and Margaret Little. Oxford: Oxford University Press, 276-304.

MacIntyre, Alasdair. 2018. *Ethics in the Conflicts of Modernity*. Cambridge: Cambridge University Press.

———. 1981. *After Virtue: A Study in Moral Theory.* South Bend: University of Notre Dame Press.

Mantel, Susan. 2018. *Determined by Reasons: A Competence Account of Acting for a Normative Reason*. London: Routledge Press.

McDowell, John. 1979. "Virtue and Reason." *The Monist*. 3: 331-350.

McKegney, Sam Sean McKeever, Ridge, Michael Ridge. 2006. "Holism About Reasons." In *Principled Ethics: Generalism as a Regulative Ideal*. Oxford: Oxford University Press.

McNaughton, David. 1988. *Moral Vision: An Introduction to Ethics*. Maiden, Massachusetts: Blackwell Publishing.

Millgram, Elijah. 2005. Ethics Done Right: Practical Reasoning as a Foundation for Moral Theory. Cambridge: Cambridge University Press.

Murdoch, Iris. 1999a. “The Sublime and the Good.” In Existentialists and Mystics: Writings on Philosophy and Literature. Edited by Peter Conradi and George Steiner. New York: Penguin Books.

———. 1999b. “The Darkness of Practical Reason.” In Existentialists and Mystics: Writings on Philosophy and Literature. Edited by Peter Conradi and George Steiner. New York: Penguin Books.

———. 1999c. “Metaphysics and Ethics.” In Existentialists and Mystics: Writings on Philosophy and Literature. Edited by Peter Conradi and George Steiner. New York: Penguin Books.

———. 1999d. “Vision and Choice in Morality.” In Existentialists and Mystics: Writings on Philosophy and Literature. Edited by Peter Conradi and George Steiner. New York: Penguin Books.

———. 1999e. "On 'God' and 'Good'". In Existentialists and Mystics: Writings on Philosophy and Literature. Edited by Peter Conradi and George Steiner. New York: Penguin Books.

———. 1970. The Sovereignty of Good. New York: Shocken Books.

Nagel, Thomas. 1979. *Mortal Questions*. Cambridge: Cambridge University Press.

———. 1986. *The View From Nowhere*. Oxford: Oxford University Press.

Nauckhoff, Josefine. 2003. "Incentives and Interests in Kant's Moral Psychology." *History of Philosophy Quarterly* 20 (1): 41-60.

Nguyen, C. Thi. 2019. "Monuments as commitments: How Art Speaks to Groups and How Groups Think in Art." *Pacific Philosophical Quarterly* 100 (4): 971-994.

Nussbaum, Martha. 2003. "Rawls and Feminism." In *The Cambridge Companion to Rawls*. Edited by Samuel Freeman. Cambridge, UK: Cambridge University Press.

Panizza, Silvia. 2019. "Moral Perception Beyond Supervenience: Iris Murdoch's Radical Perspective." *The Journal of Value Inquiry*. 53, 1-16.

Parfit, Derek. 1997. "Reasons and Motivation." *Proceedings of the Aristotelian Society (Supplementary Volume)* 71: 99–129.

Paul, L.A. 2014. *Transformative Experience*. Oxford: Oxford University Press.

Prior, A.N. 1949. "Determinables, Determinants, and Determinants (I)." *Mind* 58 (229): 1-20.

Railton, Peter. 1984. "Alienation, Consequentialism, and the Demands of Morality." *Philosophy and Public Affairs* 13 (2): 134-171.

Rawls, John. 1993. *Political Liberalism*. New York: Columbia University Press.

———. 1980. "Kantian Constructivism in Moral Theory" in *Journal of Philosophy* 77 (9): 515-572.

Raz, Joseph. 2000. "The Truth in Particularism." In *Moral Particularism*. Edited by Brad Hooker and Margaret Little. Oxford: Oxford University Press, 48-78.

Scanlon, T.M. 2017. "Normative realism and ontology: reply to Clarke-Doane, Rosen, and Enoch and McPherson." *Canadian Journal of Philosophy* 47 (6):877-897.

———. 2014. *Being Realistic About Reasons*. Oxford: Oxford University Press.

———. 1999. *Practical Reasoning and Norms*. London: Hutchinson & Company.

———. 1998. *What We Owe to Each Other*. Cambridge: Harvard University Press.

Scheffler, Samuel. 1994. "The Rejection of Consequentialism: *A Philosophical Investigation of the Considerations Underlying Rival Moral Conceptions*. Revised Edition. Oxford: Clarendon.

———. 1992. *Human Morality*. Oxford: Oxford University Press.

Schroeder, Mark. 2021. *Reasons First*. Oxford: Oxford University Press.

———. 2007. *Slaves of the Passions*. Oxford: Oxford University Press.

Schwartzman, Lisa H. 2006. *Challenging Liberalism: Feminism as Political Critique*. University Park, PA: Pennsylvania State University Press.

Setiya, Kieran. 2013. "Murdoch on the Sovereignty of Good." *Philosophers Imprint* 13: 1–21.

Sidgwick, Henry. 1981. *The Methods of Ethics*. Indianapolis: Hackett Publishing Company.

Silverman, Eric. 2019. *The Supremacy of Love: An Agape-Centered Vision of Aristotelian Virtue Ethics*. Washington, D.C.: Lexington Books.

Singer, Peter. 2011. *Practical Ethics*. Cambridge: Cambridge University Press.

Sinnott-Armstrong, Walter. 1999. "Some Varieties of Particularism." *Metaphilosophy* 30 (January-April): 1-12.

Smith, Michael. 1994. *The Moral Problem*. Oxford: Blackwell Publishing.

Stocker, Michael. 1976."The Schizophrenia of Modern Ethical Theories." *The Journal of Philosophy* 73 (14): 453-466.

Stoutland, Frederick. 2001. "Responsive action and the belief-desire model." *Grazer Philosophische Studien* 61 (1): 83-106.

Stratton‐Lake, Philip. 2000. *Kant, Duty, and Moral Worth*. New York: Routledge.

Tappolet, Christine. 2004. “Through Thick and Thin: Good and Its Determinates.” *Dialectica* 58 (2): 207–221.

Tsu, Peter. 2018. “Reason Holism, Individuation and Embeddedness.” *Ethical Theory and Moral Practice.* 21 (5): 1091-1103.

Turri, John. 2011. "Believing For a Reason." *Erkenntnis* 74 (3): 383-397.

Velleman, David. 1999. "‘Love as a Moral Emotion." *Ethics*. 109: 338–374.

Wedgwood, Ralph. 2006. "The Normative Force of Reasoning." *Noûs* 40 (4): 660–686.

Weil, Simone. 2009. *Waiting for God*. Translated by Emma Craufurd. New York: HarperCollins.

Williams, Bernard. 1985. *Ethics and the Limits of Philosophy*. Cambridge, MA: Harvard University Press.

———. 1981a. "Internal and External Reasons." In *Moral Luck*: Philosophical Papers 1973-1980. Cambridge: Cambridge University Press.

———. 1981b. "Persons, Character and Morality." In *Moral Luck*: Philosophical Papers 1973-1980. Cambridge: Cambridge University Press.

———. 1973. "A Critique of Utilitarianism." In *Utilitarianism: For and Against*. By J.C.C. Smart and Bernard Williams. Cambridge: Cambridge University Press.

Wolf, Susan. 2012. "“One Thought Too Many”: Love, Morality, and the Ordering of Commitment." In *Luck, Value, and Commitment: Themes From the Ethics of Bernard Williams*. Edited by Ulrike Heuer and Gerald Lang. Oxford: Oxford University Press.

———. 1982. "Moral Saints." *Journal of Philosophy*. 79 (8) 419-439.