

UC Riverside

UC Riverside Electronic Theses and Dissertations

Title

Taking Each Other Seriously: Towards a Theory of Morally Responsible Agency

Permalink

<https://escholarship.org/uc/item/7m64x6sj>

Author

Beglin, David P.

Publication Date

2018

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA
RIVERSIDE

Taking Each Other Seriously:
Towards a Theory of Morally Responsible Agency

A Dissertation submitted in partial satisfaction
of the requirements for the degree of

Doctor of Philosophy

in

Philosophy

by

David P. Beglin

June 2018

Dissertation Committee:

Dr. Agnieszka Jaworska, Co-Chairperson

Dr. Coleen Macnamara, Co-Chairperson

Dr. John Martin Fischer

Dr. Dana Nelkin

Dr. Andrews Reath

Copyright by
David P. Beglin
2018

The Dissertation of David P. Beglin is approved:

Committee Co-Chairperson

Committee Co-Chairperson

University of California, Riverside

ACKNOWLEDGMENTS

A friend of mine, upon finishing her dissertation, expressed anxiety about writing her acknowledgments. At the time, I didn't appreciate the reason for her anxiety. Now, though, I do. There is simply no way to express how grateful I am to everyone who has helped me and has been there for me over the years. Here, I'll do my best; although as is often the case, my words won't do my feelings justice.

UCR has been a special place to study philosophy. I've greatly enjoyed my time here, surrounded by encouraging, friendly, and thoughtful peers and professors, and I've learned much from the various seminars, reading groups, workshops, and colloquia that I've had the privilege to take part in. I should single out three groups in particular. As a first year, I found a community in the agency workshop, a community that continued to provide ideas, criticism, fun, and encouragement throughout my time at UCR. Many thanks to the members of that workshop, and especially to Zac Bachman, Taylor Cyr, Andrew Law, Jonah Nagashima, and Meredith McFadden. Later, in my second year, I began attending the Wittgenstein reading group, encouraged by Eric Walker and my experience as a TA for Phil 7, with Larry Wright. This was auspicious. I learned so much about doing philosophy with this group. I'm especially grateful to Larry, Eric, Bill Bracken, Max Murphy, Patrick Ryan, and Howie Wettstein for all of the great discussion. And finally, I've happily been adopted by a group of John Martin Fischer's past students. Of them, Neal Tognazzini has been particularly encouraging and supportive—many thanks are due to him.

I've had the pleasure of meeting some great friends at UCR. Here, for the sake of space (and respect for the limits of language), I won't try to say what they mean to me; but rest assured, they mean a lot. So, thanks to Zac Bachman, Maude Clark, Meredith McFadden, Avery Snelson, Will Swanson, Eric Walker, and Monique Wonderly. I should also thank my friends and family back home—my friends for providing endless distraction and good times; my family for the support and encouragement (and good times, too!).

I'd also, of course, like to thank my committee. Everyone on it has contributed so much to my understanding of the philosophical issues that I've been grappling with. Andy Reath may not realize it, but this project actually started in his seminars during my first year. His seminars on value theory and Scanlon were hugely influential on my trajectory as a philosopher. John Fischer has been an encouraging and thoughtful mentor throughout my time at UCR, from the very start, and I'm deeply grateful for that. And Dana Nelkin has been so incredibly generous, with both her time and her incisive comments. I don't know how she does it, but I'm thankful she does! Finally, Coleen Macnamara and Agnieszka Jaworska were the best advisors I could have hoped for. I'm so incredibly thankful for all of the time and energy they've spent reading and discussing my work with me, and for their encouragement and investment, along with their sensitivity to the ups and downs of grad school. Perhaps the thing I'm proudest of is being their student.

Finally, I should acknowledge University of Chicago Press, for allowing me to reprint (in chapter one) a couple of pages from a previously published paper: David

Beglin, "Responsibility, Libertarians, and the 'Facts as We Know Them': A Concern-Based Construal of Strawson's Reversal," *Ethics* 128 (April 2018): 612-625. Copyright 2018 by University of Chicago Press.

For my mom
(She might prefer “mum”)

ABSTRACT OF THE DISSERTATION

Taking Each Other Seriously:
Towards a Theory of Morally Responsible Agency

by

David P. Beglin

Doctor of Philosophy, Graduate Program in Philosophy
University of California, Riverside, June 2018
Dr. Agnieszka Jaworska and Dr. Coleen Macnamara, Co-Chairpersons

Regarding one another as responsible, and responding to one another accordingly, is an essential part of social and moral life. We resent each other, we feel indignation and gratitude, we have our feelings hurt, and we find ways to forgive. Still, we don't regard everyone as responsible. Some agents—young children, for example, or people suffering from certain forms of mental illness—aren't thought to be appropriate objects of such attitudes, aren't taken to be appropriately held responsible for what they do. What accounts for this? What does it mean to be a morally responsible agent, to be an appropriate object of the various attitudes that characterize our responsibility practices?

In this dissertation, I develop a distinctive approach to accounting for morally responsible agency. I develop my approach in contrast with theorists who, inspired by P.F. Strawson, understand what it means to be a morally responsible agent in terms of the

nature of the attitudes that characterize our responsibility practices, the so-called “reactive attitudes.” I argue that we ought to understand responsible agency not in terms of these attitudes themselves but, rather, in terms of the kind of concern that leaves us prone to them in the first place. I call this the “basic concern,” and when we have it towards someone, I hold, we take that person's attitudes to carry a kind of authority, to be capable of challenging or affirming our own perspective about what attitudes we should have or that we should perceive as appropriate in the relevant context. Whether it is appropriate to regard someone as a responsible agent, I argue, depends on whether it is appropriate to have this concern about that person's attitudes. A responsible agent, on my view, is thus a kind of fellow participant in social and moral life.

My account of the basic concern illuminates what's at stake in our responsibility practices, and it thus illuminates the distinctive significance that we ascribe to morally responsible agents' actions and attitudes. It also, I suggest, sheds light on the nature of the various attitudes that the basic concern leaves us prone to feel.

Table of Contents

Introduction	1
1. A Concern-Based Construal of Strawson's Reversal	9
2. Two Strawsonian Strategies	46
3. The Basic Concern	79
4. Towards a Concern-Based Theory of Morally Responsible Agency	136
Conclusion	163
<i>Bibliography</i>	166

Introduction

Regarding one another as responsible, and responding to one another accordingly, is an essential part of the various forms of community and relationship that characterize social and moral life. When a colleague dominates a meeting, excluding others' voices, we might feel indignant, perhaps confronting him or later complaining to a friend about his toxic behavior; when an acquaintance makes a dismissive remark, we might distance ourselves from her, even come to resent her; and when a stranger lends us a helping hand, we might swell with gratitude or approbation. Others respond to us similarly. These kinds of reactions are part-and-parcel with sharing social space; they're the currency of social and moral life, so to speak.

Still, we don't respond to everyone in these ways. There are some agents—children, for instance, or people suffering from certain forms of mental illness—whom we view, to varying degrees, as outside the bounds of social and moral life. While we might accord such agents rights, while we might recognize the importance of their interests, we don't accord their actions and attitudes the same status as we do a typical adult's; we don't regard them as morally responsible agents.¹

What accounts for this, though? What does it mean to be a morally responsible agent? As Susan Wolf points out, we seem to ascribe a kind of “significance” or “depth”

¹ Throughout this dissertation, I will use the phrases “moral responsibility,” “morally responsible,” and “morally responsible agent.” I don't mean to build much into the use of “moral” here. I simply mean to contrast responsibility in a theoretical or causal sense from the sense of responsibility at stake in our ordinary social and moral practices. In this respect, we can be morally responsible for actions or attitudes that might not, strictly speaking, fall within the moral domain. I take it that this fits with the way theorists commonly use “moral responsibility.” Recently, Susan Wolf has suggested we call this latter kind of responsibility “deep responsibility” (Wolf 2015).

or “power” to such agents; “we take responsible beings more seriously” in some way (Wolf 1990, 5-6). But while these phrases are suggestive, it remains unclear how we should understand them. What is the nature of the significance that we ascribe to a morally responsible agent? And what sorts of capacities are required for one’s actions and attitudes to be appropriately taken to carry this significance, for one to be appropriately regarded as a morally responsible agent?

In this dissertation, I develop a distinctive approach to addressing these questions, to accounting for morally responsible agency. Central to my approach is what I call, following Gary Watson (2014), the “basic concern.” This is the concern that leaves us prone to hold people responsible, the concern around which our responsibility practices revolve. The concern-based approach that I advocate understands morally responsible agency in terms of the nature of this basic concern.

I develop this concern-based approach in conversation with theorists who draw inspiration from P.F. Strawson’s landmark (1962) essay, “Freedom and Resentment.” “Freedom and Resentment” is well known for a number of reasons, but particularly relevant for my purposes are two contributions that it made to the literature on moral responsibility. First, Strawson there introduced the class of “reactive attitudes.” These are the various attitudes that characterize our practices of holding responsible, including, *inter alia*, resentment, guilt, indignation, gratitude, hurt feelings, and forgiveness. Many theorists, even those who distance themselves from other aspects of Strawson’s thought, have come to understand holding responsible in terms of these attitudes. But Strawson introduced the reactive attitudes in order to make a larger point about responsibility.

“Only by attending to this range of attitudes,” he writes, “can we recover from the facts as we know them a sense of what we mean, i.e. of *all* we mean, when, speaking the language of morals, we speak of desert, responsibility, guilt, condemnation, and justice” (Strawson 1962, 64; his emphasis). This marks the second major contribution that Strawson made to the literature on moral responsibility. Strawson held that what it means to be morally responsible (and a morally responsible agent) depends, in some way, on our practices of holding responsible. This is often called his “reversal” thesis (Watson 1987, 257-259).

There is no real consensus about how, precisely, we should understand Strawson’s reversal. If anything, there remains a great deal of confusion about this aspect of “Freedom and Resentment” (see Todd 2016 for discussion). Most interpretations of the reversal, though, emphasize the reactive attitudes. In my first chapter, I depart from this orthodoxy. I situate “Freedom and Resentment” within Strawson’s larger corpus, particularly in relation to his work on the nature of philosophical analysis. I argue that we should understand “Freedom and Resentment” as an exercise in Strawson’s preferred mode of philosophy, what he called “connective analysis” (see Strawson 1992). We should thus see Strawson as attempting to elucidate the concepts involved in our responsible practices by attending to their use and, more importantly, by offering an explanation of why we have such practices in the first place. This latter aspect of connective analysis, understanding responsibility in terms of the core structural feature of our responsibility practices, is what leads me to my concern-based construal of Strawson’s reversal. As I understand Strawson, he takes what it means to be morally

responsible to depend on the concern that gives rise to our responsibility practices, that leaves us prone to the reactive attitudes that characterize those practices—the basic concern.

The concern-based construal of Strawson’s reversal suggests that we ought to account for morally responsible agency in terms of the nature of the basic concern, which constitutes the evaluative stance from which we hold people responsible. I call this way of accounting for morally responsible agency the “concern-based Strawsonian strategy.” This contrasts with the predominant way that theorists have taken up and developed Strawson’s thought.

On the predominant Strawsonian approach, theorists account for morally responsible agency in terms of the nature of the reactive attitudes, the way those attitudes engage their objects. For example, many Strawsonian theorists have argued that the reactive attitudes morally address their object, communicating a demand or message. For these theorists, a morally responsible agent is an intelligible object of such moral address (Watson 1987, 2011; Darwall 2006; Shoemaker 2007; Macnamara 2015b). Other theorists, most prominently R. Jay Wallace (1994), have argued that the reactive attitudes, or, at least, resentment, indignation, and guilt, essentially hold people to obligations. They’ve thus suggested that a morally responsible agent is the type of agent who can reasonably be held to such obligations. And still other theorists have argued that the reactive attitudes are essentially responses to the quality of will expressed in one’s actions. A morally responsible agent, these theorists hold, is the kind of agent whose

actions can express the relevant kind of good or ill will (McKenna 2012, Talbert 2014, Shoemaker 2015, 2017).

While Strawsonian theorists have developed a wide variety of views about morally responsible agency, then, they've shared a common strategy. They account for morally responsible agency in terms of what the reactive attitudes are, in terms of the nature of these attitudes, the features that they have.

In my second chapter, I argue that the concern-based Strawsonian strategy is more promising than this predominant "attitude-based Strawsonian strategy." First, the attitude-based approach faces problems that the concern-based approach avoids. The reactive attitudes, I argue, are too various and too variable to provide a stable foundation for a theory of morally responsible agency. We'd do better, I hold, to focus on something less superficial than those attitudes, namely, the concern that gives rise to them in the first place. Moreover, the concern-based approach seems better situated to account for the distinctive significance that we ascribe to someone when we regard that individual as a morally responsible agent. Whereas the reactive attitudes reveal *how* we engage with morally responsible agents, the basic concern reveals *why* we engage with morally responsible agents in that way. The basic concern thus sheds more direct light on the significance that a morally responsible agent's actions and attitudes have for us.

How, though, should we understand the basic concern? I turn to this question in my third chapter. Theorists haven't said much about the basic concern, no doubt due to the tendency to focus instead on the reactive attitudes. I thus take my third chapter to represent a significant contribution to the moral psychology of responsibility.

I begin by distinguishing between two aspects of the basic concern. First, there is a question about the concern's object, what it is a concern *about*. This is the aspect of the basic concern to which theorists have given the most attention. Indeed, Strawson, in "Freedom and Resentment," understood the basic concern to be a concern about the quality of will with which people act, and many theorists have developed views that have drawn on this aspect of Strawson's thought. But, as I point out in my first chapter, Strawson overlooked the fact that we can care about the same object in different ways. This constitutes a major lacuna in his theory, and, I argue, explains the difficulty he has accounting for morally responsible agency.² The second aspect of the basic concern, then, has to do with the nature or mode of the basic concern, the *way* we care when we have it.

On my account, the basic concern is a concern about how someone regards the relevance of considerations that we think should matter in the social contexts that we share with that person. When we have this concern, I hold, we take that person's regard, their evaluative outlook, to be capable of making a claim on us about how we should regard the relevance of those considerations. In this respect, we see that person as a fellow participant with us, someone whose evaluative outlook carries a kind of authority, whose outlook can challenge or affirm our own perspective about what is appropriate in the relevant context.

This account of the basic concern, I argue, has a number of virtues. First, it provides a way of vindicating Strawson's claims that the basic concern is a ubiquitous

² For important criticisms of Strawson's account of morally responsible agency, see Bennett 1980, Watson 1987, Russell 1992, and Wallace 1994. It's worth remarking that all of these theorists are sympathetic to Strawson's project.

and deep-seated feature of human psychology. Indeed, I suggest, this concern is reflective of the way we develop into social and moral agents in the first place. Second, this account of the basic concern sheds light on the sense in which regarding someone as a morally responsible agent is a kind of respect: it gives authority to their perspective in the social contexts that we share with them. And, finally, the basic concern can plausibly make sense of the various reactive attitudes that Strawson adduces, including resentment, indignation, gratitude, forgiveness, guilt, hurt feelings, and certain kinds of love. This latter virtue is particularly important for the moral psychology of responsibility. Strawsonian theorists have struggled to account for the wide-class of reactive attitudes; they often narrow the reactive attitudes to a small set, usually resentment, indignation, and guilt. My account of the basic concern, on the other hand, seems to offer a promising way to understand the reactive attitudes as a class, including attitudes that are frequently neglected by theorists who employ the attitude-based Strawsonian strategy, attitudes like hurt feelings and forgiveness.

With an account of the basic concern on the table, I turn back to the topic of morally responsible agency in my fourth and final chapter. In this chapter, which is my most speculative, I explore some of the possible implications that my account of the basic concern might have for how we should understand morally responsible agency, sketching the starting points for a theory about it. I suggest that a morally responsible agent is a kind of fellow participant with us, someone whose evaluative outlook carries authority, is capable of making a claim on us about how we should regard what matters in the context we share with them. To be appropriately regarded as such an agent, I argue, seems to

require one have at least three capacities. First, it seems one must be capable of acting and forming attitudes under the guise of the good. In this respect, one must be capable of committing to one's own evaluative outlook as appropriate. Second, it seems like one must have normative competence, which includes having both a general understanding of what matters in the context one is in and the executive capacity to bring that understanding to bear on one's practical reasoning in that context. And, finally, it seems that one must have moral sense; one must be capable of conceiving of oneself as a fellow participant among others. All three of these capacities, I argue, appear to be presumed when we see someone as a fellow participant, when we have the basic concern about their evaluative outlook.

Chapter 1

A Concern-Based Construal of Strawson's Reversal

1. Introduction

Only a handful of works have had as large of an influence on the philosophical literature on moral responsibility as P.F. Strawson's "Freedom and Resentment." And of these influential works, few have remained as mysterious. My aim in this chapter is to shed some light on one of Strawson's key ideas in that essay, his idea that what it means to be morally responsible depends (in some way) on our practices of holding responsible—his so-called "reversal" thesis (see Watson 1987, 121-122).

Strawson's reversal plays a key role in "Freedom and Resentment." It's also a central aspect of the more general, contemporary Strawsonian program. Numerous theorists, inspired by Strawson, have developed theories of morally responsible agency by drawing on something like it. In understanding Strawson's reversal, and in developing it, theorists have typically emphasized the reactive attitudes that undergird our responsibility practices, attitudes like resentment, gratitude, indignation, and so on. In the next chapter, for instance, I'll discuss and criticize what is perhaps the most predominant Strawsonian approach, which understands morally responsible agency in terms of these attitudes' *nature*.¹ Before turning to this "attitude-based" Strawsonian approach, though, I'll offer, in this chapter, an interpretation of Strawson that deemphasizes the reactive attitudes. On my interpretation, Strawson's idea is that what it means to be morally

¹ For theorists who take this approach, see, e.g., Watson 1987, 2011; Wallace 1994; Darwall 2006; Shoemaker 2007, 2011, 2015, 2017; McKenna 2012; Talbert 2012; and Macnamara 2015b.

responsible depends on the concern that leaves us susceptible to the reactive attitudes, the concern around which our responsibility practices revolve. In this respect, I hold, the reactive attitudes themselves aren't as central to understanding Strawson as many have taken them to be.²

To develop my interpretation, I'll draw on Strawson's larger corpus, particularly his work on the nature of philosophical analysis. This is an area of Strawson's thought that has largely been neglected in attempts to understand "Freedom and Resentment," much, I think, to the detriment of the literature on responsibility.³ On my interpretation, "Freedom and Resentment" is an instance of Strawson's preferred mode of analysis, *connective analysis*. Better understanding connective analysis thus puts us in a better position to understand the strategy that Strawson pursues in "Freedom and Resentment." It also puts us in a better position to understand how Strawson errs in "Freedom and Resentment," or, at least, to understand the limits of his discussion there, a fact that will help set the stage for the larger project of this dissertation.

2. The Interpretive Problem

Before considering Strawson's work on the nature of philosophical analysis, we should first briefly turn to "Freedom and Resentment" itself. This will allow us to locate the particular interpretive problems that I'd like to address.

² To be clear, this isn't to say that the reactive attitudes aren't important at all, or that the introduction of the reactive attitudes doesn't represent a significant contribution to our understanding of responsibility. The point, rather, is simply that the reactive attitudes aren't *as* central as theorists have taken them to be.

³ Recently, theorists have begun to devote more attention to Strawson's surrounding work. See, e.g., Campbell 2017, Coates 2017, and Pamela Hieronymi's unpublished manuscript, "Freedom, Resentment, and the Metaphysics of Morals."

In “Freedom and Resentment,” Strawson is engaged with a dispute concerning the compatibility of determinism and moral responsibility. This dispute is between two disputants, the optimist and the pessimist, and it has the following shape. The optimist holds that moral responsibility is compatible with determinism. This is because she takes our responsibility practices to be based on their efficacy in “regulating behaviour in socially desirable ways” (Strawson 1962, 46). Because determinism in no way changes the fact that holding someone responsible can bring about socially desirable effects, according to the optimist, whether the thesis of determinism is true is irrelevant to our responsibility practices. This consequentialist view is anathema to the pessimist, however, who insists it leaves something “vital” out of our practices. Holding someone responsible, the pessimist contends—blaming or praising someone—is only justified if the person held responsible *deserves* it, and the optimist’s view doesn’t seem able to account for this.⁴ Someone can deserve blame or praise, according to the pessimist, only if that person is free in a strong sense, a sense incompatible with determinism. So, the pessimist concludes, moral responsibility must be incompatible with determinism. Strawson’s aim in “Freedom and Resentment” is to reconcile the optimist and pessimist, offering an account of the basis of our responsibility practices that supports a

⁴ It’s worth noting that we need not construe the pessimist’s point as a point about retribution. Although desert is often understood retributively, there are non-retributive construals of it. We might talk in terms of praise- and blameworthiness, for instance, or of praise and blame being fitting (see Scanlon 2008, 188-189). The pessimists’ worry isn’t that the optimist can’t account for a sense in which wrongdoers should suffer.

compatibilist conclusion, like the optimist's view, but that can also capture the sense of desert the pessimist is rightly worried about preserving.⁵

To accomplish this aim, Strawson draws his audience's attention to the "reactive attitudes," the various emotions and feelings involved in our responsibility practices. These include "personal," "non-detached" attitudes, like "gratitude, resentment, forgiveness, love, and hurt feelings"; "vicarious" attitudes, which can be felt on behalf of another person, such as indignation and moral approbation; and attitudes that we have towards ourselves, like guilt, the feeling of obligation, and shame (see Strawson 48-50, 56-58). These three clusters of reactive attitudes—personal, vicarious, and self-directed—are "humanly" connected, according to Strawson; they have, he explains, "common roots in our human nature and our membership in human communities" (58).

For Strawson, this range of attitudes is in some way the key to understanding the concepts that are central to our practices of responsibility. He thus writes, in a passage that addresses the dispute between the optimist and the pessimist:

The vital thing can be restored by attending to that complicated web of attitudes and feelings which form an essential part of the moral life as we know it, and which are quite opposed to objectivity of attitude. Only by attending to this range of attitudes can we recover from the facts as we know them a sense of what we mean, i.e. of *all* we mean, when, speaking the language of morals, we speak of desert, responsibility, guilt, condemnation, and justice. (64; his emphasis)

This passage is clearly crucial to understanding Strawson's project. Strawson, after all, is ultimately concerned with (1) showing the pessimist that the foregoing concepts don't implicate determinism in the way the pessimist thinks they do and (2) showing the

⁵ For Strawson's account of the dialectic between optimists and pessimists, along with his aims with regard to it, see Strawson 1962, 45-48.

optimist that there is a better way to account for these concepts than a consequentialist foundation of our responsibility practices. To fulfill these aims, Strawson evidently needs to proffer his own foundation for these concepts, one that shows that they don't implicate determinism. And here, in this passage, we get a glimpse of that foundation: we must attend "to that complicated web of attitudes and feelings which form an essential part of the moral life as we know it, and which are quite opposed to objectivity of attitude."

The foregoing passage, then, is crucial for understanding Strawson's strategy in "Freedom and Resentment," and particularly for understanding how he hopes to reconcile the optimist and pessimist. It also gestures at a thesis that is philosophically interesting in its own right. Somehow, Strawson is suggesting, concepts like moral responsibility are grounded in our social sentimental nature; we can only understand these concepts in terms of a range of attitudes to which we're naturally prone.⁶

Still, how precisely is this supposed to work? What exactly is Strawson suggesting? It's notoriously difficult to say. As Patrick Todd (2016) has recently pointed out, after meticulously canvassing extant articulations of Strawson's reversal, theorists have generally attributed one of two views to Strawson, both of which face problems.

Some theorists have suggested that Strawson holds that what it means to be morally responsible is simply a matter of our patterns of reactive response. On this interpretation, someone's being morally responsible is determined by whether we're disposed to respond to that person with the reactive attitudes. Such a view, though, leaves no room for criticism of our responsibility practices, no room for criticism of our

⁶ I borrow the phrase "social sentimental nature" from Watson 2014.

dispositions to respond to certain people with the reactive attitudes. Thus, for example, if a community were disposed to hold young children responsible, to respond to them with the reactive attitudes, then those young children would *be* responsible. And this seems problematic. It doesn't allow for the possibility that the community can get it wrong.⁷

The alternative view avoids this problem by adding normative language. On this alternative, Strawson takes someone's being morally responsible to depend on whether it is *appropriate* (to be disposed) to respond to that person with the reactive attitudes. This interpretation allows for the necessary critical distance between whether someone is morally responsible and whether we're disposed to hold that person responsible; however, it does so at a significant cost. This latter interpretation leaves unexplained what we're ultimately interested in understanding: namely, the relevant sense in which and conditions under which it is "appropriate" to be disposed to respond to someone with the reactive attitudes (cf. Nelkin 2011, 29). In this respect, this latter interpretation leaves Strawson unable to address the dispute between optimists and pessimists, unable to explain why determinism doesn't bear on moral responsibility. After all, determinism might make it inappropriate to be disposed to respond to people with the reactive attitudes (Todd 2016).

Recently, David Shoemaker (2017) has attributed a more complex "response-dependent" view to Strawson. This view analyzes moral responsibility in terms of the fittingness of the reactive attitudes. Here, "fittingness" gestures at a particular kind of normativity, one that inheres in the attitudes themselves. That is, an attitude is fitting just

⁷ For this criticism, see, e.g., Fischer 1994, 212-213; Ekstrom 2000, 149; Nelkin 2011, 28; and Todd 2016.

in case it accurately construes the world. Shoemaker unpacks the idea in terms of blameworthiness and anger. Someone is blameworthy, on this response-dependent view, just in case what they did merits anger, and whether someone merits anger depends on the nature of anger, particularly on what anger essentially responds to.

Shoemaker's "fitting response-dependent" view is a reflection of a broader Strawsonian approach, which I'll discuss and criticize in the next chapter. On this broader "attitude-based" Strawsonian approach, one accounts for morally responsible agency in terms of the *nature* of the reactive attitudes. I believe that there are philosophical reasons to worry about this approach.

Still, here I'm interested in how we should understand Strawson, and whether the attitude-based approach is philosophically problematic is a different question from whether Strawson took that approach. I doubt, though, that he did. In the section where he introduces the reactive attitudes, for example, Strawson seems more interested in "the general causes of the reactive attitudes" than the nature of the reactive attitudes themselves (Strawson 1962, 50). "The central commonplace that I want to insist on," Strawson writes, "is the very great importance that we attach to the attitudes and intentions towards us of other human beings, and the great extent to which our personal feelings and reactions depend upon, or involve, our beliefs about these attitudes and intentions" (48). Of course, it might seem like Strawson is getting at something here that is similar to what Shoemaker has in mind. It might seem, that is, like Strawson is suggesting that we need to focus on what the reactive attitudes respond to. But I suspect that things are more complicated. Strawson goes on to write that what he wants to

emphasize is “how much we actually mind, how much it matters to us, whether the actions of other people—and particularly of *some* other people—reflect attitudes of goodwill, affection, or esteem on the one hand or contempt, indifference, or malevolence on the other” (49). For Shoemaker, whether someone is morally responsible depends on whether their actions are capable of reflecting the relevant quality of will to which the reactive attitudes respond.⁸ In this last quote, though, Strawson seems to move away from this idea. What Strawson wants to emphasize is the way that we care about *some* people’s quality of will. And indeed, Strawson later suggests that there are some agents whom we don’t consider morally responsible, despite the fact that their actions can reflect contempt, indifference, and malevolence—the sorts of attitudes and intentions to which we might normally respond with anger or resentment (see Strawson 1962, 51-53).

Of course, at this point I don’t mean to suggest that I’ve settled anything. My point, rather, is that there is room here to wonder about what precisely Strawson is drawing our attention to when he tells us to attend “to that complicated web of attitudes and feelings which form an essential part of the moral life as we know it.” Moreover, we might wonder about *why* we must attend to this “web” or “range” of attitudes. What lies behind Strawson’s general suggestion? What connection does he see between this “range of attitudes” and the notion of moral responsibility?

I think that the foregoing questions are interrelated and that to get a grip on them we need to take a step away from “Freedom and Resentment.” What we need to understand is Strawson’s general philosophical methodology. Fortunately, Strawson had

⁸ This is a common Strawsonian view. See, e.g., McKenna 2012, Talbert 2012, and Shoemaker 2015.

much to say on this topic. Let's leave "Freedom and Resentment" for a moment, then, and try to get a broader view of Strawson's approach.

3. Connective Analysis and Descriptive Metaphysics

In *Analysis and Metaphysics*, Strawson introduces a new conception of analysis, a new model of philosophical inquiry, which he calls "connective analysis."⁹ He contrasts this way of doing philosophy with what he sees as the more typical style: reductive analysis. Reductive analysis aims to find ideas that are "completely simple"—"free from internal conceptual complexity"—and to "demonstrate how the more or less complex ideas that are of interest to philosophers could be assembled by a kind of logical or conceptual construction out of these simpler elements" (Strawson 1992, 17). Thus, for example, we might reduce knowledge to justified true belief. Connective analysis is meant as an alternative to this reductive model of philosophical inquiry. Strawson explains:

Let us abandon the notion of perfect simplicity in concepts; let us abandon even the notion that analysis must always be in the direction of greater simplicity. Let us imagine, instead, the model of an elaborate network, a system of connected items, concepts, such that the function of each item, each concept, could, from a philosophical point of view, be properly understood only by grasping its

⁹ *Analysis and Metaphysics* was published in 1992. This is thirty years after "Freedom and Resentment." One might thus worry that my using *Analysis and Metaphysics* in this context is anachronistic. But though it was published in 1992, *Analysis and Metaphysics* is a published version of Strawson's "introductory lectures in philosophy," which he gave (nearly) annually from 1968 to 1987. Here, I'm drawing on the first two chapters of the book, and Strawson remarks in his preface that these chapters, and the five that follow them, "preserve, virtually unchanged, the content of the first ten or eleven of the 1968 lectures" (vii). The chapters I'm drawing from, then, were written in 1968, only 6 years after "Freedom and Resentment." Moreover, as will come out in my discussion, the ideas Strawson discusses in *Analysis and Metaphysics* are connected to ideas Strawson wrote about from 1953 to as late as 1985.

connections with the others, its place in a system—perhaps better still, the picture of a set of interlocking systems of such a kind. (Strawson 1992, 19)

Concepts are located in systems, then; they're connected to each other in a systematic way. And to grasp their significance, to get a sense of what concepts mean, we have to understand them in terms of these systems.

How does this systematicity come about? Strawson presents a picture of human beings participating in practices. These are shared forms of life, so to speak, systems of rules or principles that guide a particular activity, and in which people participate with a tacit mastery of how to go on together. Our concepts find their home in our practices. And as participants in practices, we develop rich conceptual apparatuses and learn how to apply them. Importantly, though, participating in a practice is different from being able to articulate the “theory” underlying it. “In our transactions with each other and the world,” Strawson explains, “we operate with an enormously rich, complicated, and refined conceptual equipment; but we are not, and indeed could not be, taught the mastery of the items of this formidable equipment by being taught the *theory* of their employment” (Strawson 1992, 6). He goes on:

Thus, for example, we know, in a sense, what knowing is perfectly well long before we hear (if we ever do hear) of the Theory of Knowledge. We know, in one sense, what it is to speak the truth without perhaps suspecting that there are such things as Theories of Truth. We learn to handle the words ‘the same’, ‘real’, ‘exists’, and to handle them correctly, without being aware of the philosophical problems of Identity, Reality, and Existence. In the same way, we learn to operate with a vast and heterogeneous range of notions... Of course we learn the words which express these concepts in a variety of ways; but we learn them largely without benefit of anything which could properly be called general theoretical instruction. We are not introduced to them by being told their place in a general theory of concepts. Such instruction as we do receive is severely practical and largely by example. (Strawson 1992, 6-7)

Our competence with these concepts, then, predates any theoretical understanding of what they are or to what (philosophical) problems they give rise. And our competence develops from a young age, through our participation in the practices found in our community or society.

Theoretical work—philosophical work—thus relies on there already being such practices, on our concepts already having a sense to them, given by the practice in which they find their home. Strawson draws an analogy with grammar. Providing a grammar, a theoretical articulation of the rules and principles guiding the use of a particular language, requires the existence, already, of that language. Moreover, like our practices more generally, one can be practically competent in a language without being able to articulate the rules underlying it.

Using this grammatical analogy, Strawson offers a compelling statement of what philosophy is about:

...just as the grammarian, and especially the model modern grammarian, labours to produce a systematic account of the structure of rules which we effortlessly observe in speaking grammatically, so the philosopher labours to produce a systematic account of the general *conceptual structure* of which our daily practice shows us to have a tacit and unconscious mastery. (1992, 7)

This, then, is Strawson's broad conception of philosophy; philosophy, for Strawson, aims at "general human conceptual self-understanding" (Strawson 2011, 225). I won't defend this conception of philosophy from possible criticisms here, but I believe that it has a lot going for it. As we've seen, it is grounded in and reflects a certain picture of human practice and meaning that seems realistic. In this way, Strawson's conception of philosophy is naturalistic, grounding analysis and concepts in something recognizable,

even ordinary. Moreover, it is deeply humane. Philosophy becomes a matter of making sense of human experience, elucidating our practices by drawing our attention to the conceptual structures undergirding them.¹⁰ Finally, it allows for philosophy to play a concrete, positive role in human life. By better understanding our practices and the concepts they employ, we can better elaborate and refine them.

Philosophy, for Strawson, is thus about producing illuminating, systematic accounts of the conceptual structures underlying our practices. Per his articulation of connective analysis, this means drawing connections between the concepts those conceptual structures involve. But how do we go about doing this?

One way is by attending to the everyday use of the concepts that philosophically interest us. This has some resemblance to Wittgenstein's (1953) conception of philosophy. Both his and Wittgenstein's respective conceptions of philosophy, Strawson admits, "place great weight on the actual use of concepts... They both suggest that, somehow, the saving truth lies there, in the actual employment of concepts" (1992, 8-9). Nevertheless, Strawson argues that he importantly differs from Wittgenstein. Strawson's conception has a positive orientation towards its subject matter: "there is the suggestion of a system; of a general underlying structure to be laid bare; even of explanation."

Wittgenstein's conception of philosophy, on the other hand, at least as Strawson sees it, is largely negative. As Strawson describe it: "We are not to construct a system, but to 'assemble reminders' for a particular purpose; and this purpose is that of liberating

¹⁰ When he introduces the notion of "connective analysis," Strawson remarks that he is tempted to refer to it as "elucidation," rather than "analysis." He decides against this, though, because "analysis" "is consecrated by usage and since it has... a more comprehensive sense than [dismantling or reducing]" (Strawson 1992, 19).

ourselves from the confusions and perplexities we get into when our concepts are idling in the mind” (Strawson 1992, 9).

Despite these differences, it’s worth noting that Strawson’s conception isn’t competing with Wittgenstein’s. The two are compatible; Strawson’s conception of philosophy merely adds something Wittgenstein eschews. In an earlier essay, for example, Strawson suggests that something like Wittgenstein’s conception represents one important strand of philosophy: dealing with perplexity and paradox. “For it often happens,” Strawson explains, “that someone reflecting on a certain set of concepts finds himself driven to adopt views which seem to others paradoxical or unacceptably strange, or to have consequences which are paradoxical or unacceptably strange.” When this happens, Strawson continues, the person is “temporarily dominated by one logical mode of operation of expressions, or by one way of using language, or by one logical type or category of objects, or by one sort of explanation, or by one set of cases of the application of a given concept.” To correct the distortions created by this “temporary one-sidedness of vision,” Strawson concludes, “one must make plain the actual modes of operation of the distorted concepts or types of discourse... and, in doing this, one must, if one can, make plain the sources of the blinding obsession with the model cases” (Strawson 1963, 515).¹¹

That Strawson’s conception of philosophy is compatible with Wittgenstein’s is important to note, because there is no doubt an element of Wittgensteinian “therapy” in

¹¹ There is an interesting connection between this passage and the way Strawson describes the optimist’s style of “overintellectualizing the facts” as a “characteristically incomplete empiricism, a one-eyed utilitarianism” (see Strawson 1962, 64).

“Freedom and Resentment.” But while Strawson is certainly engaged in dealing with the purported perplexity of both the optimist and the pessimist, we shouldn’t overstate this aspect of his project. Crucially, there is more to Strawson’s conception of philosophy than this negative project of dissolving perplexity and paradox.

First, in addition to this negative project, Strawson notes that we can inquire after concepts for no other purpose than to better understand them. And doing so will involve attending to their actual use in our practices. This, according to Strawson, is a second strand of philosophy.

This second strand, though, doesn’t exhaust Strawson’s positive conception of philosophy. In fact, Strawson seems to worry that merely attending to the use of our concepts doesn’t go deep or far enough. He thus describes a third strand of philosophy. This third strand is the “attempt to explain, not just how our concepts and types of discourse operate, but why it is that we have such concepts and types of discourse as we do.” He hastens to add: “This is not a historical enquiry. It attempts to show the natural foundations of our logical, conceptual apparatus” (Strawson 1963, 515-516). This third strand is what Strawson later came to call “descriptive metaphysics.”¹² Descriptive metaphysics aims “to lay bare the most general features of our conceptual structure”

¹² Strawson introduced the notion of descriptive metaphysics in his 1959 book, *Individuals*. This book was thus published before the essay I’ve been quoting, which was published in 1963. Nevertheless, despite its later publication date, that essay, “Carnap’s Views on Construct Systems Versus Natural Languages in Analytic Philosophy,” was actually written much earlier than *Individuals*. As Strawson reports in his “Intellectual Autobiography,” his essay on Carnap was written in 1954, 5 years before *Individuals* was published (Strawson 2011b, 235).

(Strawson 1959, 9). Strawson suggests that it is only through descriptive metaphysics that we can arrive at the sort of general understanding of concepts that we seek in philosophy:

Up to a point, the reliance upon a close examination of the actual use of words is the best, and indeed the only sure, way in philosophy. But the discriminations we can make, and the connexions we can establish, in this way, are not general enough and not far-reaching enough to meet the full metaphysical demand for understanding. For when we ask how we use this or that expression, our answers, however revealing at a certain level, are apt to assume, and not to expose, those general elements of structure which the metaphysician wants revealed. The structure he seeks does not readily display itself on the surface of language, but lies submerged. (Strawson 1959, 9-10)

In addition to examining how concepts are actually used, then, we're to look for the core structural features of our practices. By identifying them we can better elucidate the various concepts those practices employ.

Take, for example, induction and our practices of belief-formation. Strawson touches on these issues, along with their relation to "connective analysis," briefly in *Naturalism and Skepticism*. They come up in a discussion of Hume. Hume, Strawson suggests, takes "the habit of induction" as one of "those structural features of our conceptual scheme, the framework features, which must be regarded as equally beyond question and beyond validation, but which offer themselves, rather, for the kind of philosophical treatment which I have suggested and which might be called 'connective analysis'" (Strawson 1985, 25). We might fill this idea out with a remark that Strawson makes about Hume earlier in his discussion:

Our inescapable natural commitment is to a general frame of belief and to a general style (the inductive) of belief-formation. But *within* that frame and style, the requirement of Reason, that our beliefs should form a consistent and coherent system, may be given full play. Thus, for example, though Hume did not think that a rational justification of induction in general was either necessary or possible, he could quite consistently proceed to frame "rules for judging of

cause and effect.” Though it is Nature which commits us to inductive belief-formation in general, it is Reason which leads us to refine and elaborate our inductive canons and procedures and, in their light, to criticize, and sometimes to reject, what in detail we find ourselves naturally inclined to believe. (Strawson 1985, 14)

The idea seems to be that humans are naturally committed to forming beliefs through induction. Accordingly, practices of belief-formation will inevitably arise among human populations, and these practices will be shaped by this “habit” of induction. This fact will shed light on concepts like justification, which get their sense from these practices of inductive belief-formation. But note, these practices are not just brute consequences of our habit of induction; we’re capable of elaborating and refining them. Our concept of justification, likewise, is capable of elaboration and refinement. Still, there will be natural contours to that concept, just as there will be natural contours to our practices of inductive belief-formation, because that concept and those practices are ultimately grounded in something about *us*, something natural and non-rational, something belonging to “that central core of human thinking which has no history” (Strawson 1959, 10).¹³

¹³ Should we worry about Strawson’s emphasis on ahistorical explanations for our practices? Someone of a Nietzschean bent might. Such a theorist might suggest that Strawson is just another “English psychologist” who lacks “the historical spirit itself” (Nietzsche 1998, 10). To understand our concepts, such a theorist might argue, we need to look at how our concepts developed throughout history. Is this a challenge to Strawson’s position? Not entirely. There may be some disagreement here, but Strawson’s view can accommodate genealogical inquiries. Strawson’s point is simply that *however* our concepts developed they will have a certain shape, and this shape will be a result of the natural concerns that human beings have. These natural concerns, though, according to Strawson, are an ahistorical matter; grounded, perhaps, in evolutionary or developmental features of human beings. This may not be entirely out of step with Nietzsche, interestingly. For an account of Nietzsche’s account of free will and responsibility that draws on Strawson, see Clark 2015.

4. The Core Structural Feature of Our Responsibility Practices

I suggest, then, that we interpret “Freedom and Resentment” as an essay in the style of connective analysis. But how does that help us with our interpretive questions?

Those questions, remember, arose in light of a particular passage:

The vital thing can be restored by attending to that complicated web of attitudes and feelings which form an essential part of the moral life as we know it, and which are quite opposed to objectivity of attitude. Only by attending to this range of attitudes can we recover from the facts as we know them a sense of what we mean, i.e. of *all* we mean, when, speaking the language of morals, we speak of desert, responsibility, guilt, condemnation, and justice. (Strawson 1962, 64; his emphasis)

I believe we should interpret Strawson here as suggesting that this “complicated web of attitudes and feelings” is a core structural feature of our responsibility practices, the kind of feature a descriptive metaphysician seeks. In this way, it has the potential to elucidate the various concepts those practices involve.

Indeed, later, in the paragraph immediately following the one in which the above quote occurs, Strawson writes:

Inside the general structure or web of human attitudes and feelings of which I have been speaking, there is endless room for modification, redirection, criticism, and justification. But questions of justification are internal to the structure or relate to modifications internal to it. The existence of the general framework of attitudes itself is something we are given with the fact of human society. As a whole, it neither calls for, nor permits, an external ‘rational’ justification. (Strawson 1962, 64)

Strawson then goes on, in a footnote, to compare this to “the question of the justification of induction”:

The human commitment to inductive belief-formation is original, natural, non-rational (not *irrational*), in no way something we choose or could give up. Yet rational criticism and reflection can refine standards and their application, supply ‘rules for judging of cause and effect’. (64, footnote 7)

Here, then, we have a direct comparison drawn between “the general structure or web of human attitudes and feelings” and our commitment to inductive belief-formation. We’ve already seen that Strawson, following Hume, takes our habit of induction to be a structural feature of our conceptual scheme, “which must be regarded as equally beyond question and beyond validation” but which is a candidate for connective analysis. Insofar as it is comparable to induction, it seems this “general structure or web of human attitudes and feelings” should be a good candidate for connective analysis as well.

Still, as I suggested above, in the second section, we might wonder about what precisely Strawson is drawing our attention to here. He talks about a “complicated web,” a “range,” a “general structure” of attitudes and feelings. What do these phrases refer to? Most Strawsonian theorists have developed Strawson’s approach in terms of the *nature* of the reactive attitudes. Should we understand Strawson himself as taking this approach, though? Or is he perhaps drawing our attention to something else? Perhaps to the patterns of our reactive responses, as others have suggested?

To see what Strawson is getting at, it is best to start with the part of “Freedom and Resentment” where he first introduces his audience to the reactive attitudes. After all, this is where he first attends to these attitudes, and so his discussion there should reveal something about what he takes to be significant about them.

Recall, Strawson begins this section of the paper by explaining that the central commonplace he wants to insist on is “the very great importance that we attach to the attitudes and intentions towards us of other people, and the great extent to which our personal feelings and reactions depend upon, or involve, our beliefs about these attitudes

and intentions” (Strawson 1962, 48) And after again emphasizing “how much we actually mind, how much it matters to us, whether the actions of other people—and particularly of *some* other people—reflect attitudes of goodwill, affection or esteem on the one hand or contempt, indifference or malevolence on the other,” Strawson ends the section with the following exhortation:

We should think of the many different kinds of relationship which we can have with other people—as sharers of a common interest; as members of the same family; as colleagues; as friends; as lovers; as chance parties to an enormous range of transactions and encounters. Then we should think, in each of these connections in turn, and in others, of the kind of importance we attach to the attitudes and intentions towards us of those who stand in these relationships to us, and of the kinds of *reactive* attitudes and feelings to which we ourselves are prone. (Strawson 1962, 49; his emphasis)

Strawsonian theorists typically put emphasis on the reactive attitudes, but the foregoing passages are interesting because they suggest Strawson himself is more interested in the relations underlying our proneness to those attitudes. Strawson seems to be drawing our attention to the way our responsibility practices, and particularly the various reactive attitudes they involve, revolve around a deep-seated social concern that we have about other people’s, or, at least, *some* other people’s, attitudes and intentions. It is this concern that underpins our susceptibility to the reactive attitudes, and it is a concern that is somehow tied up with our participation in various relationships with other people. We can call it, following Gary Watson, the “basic concern” (Watson 2014, 17).

What I’m suggesting, then, is that this basic concern represents a core structural feature of our responsibility practices, for Strawson. But if this is the case, why does Strawson tell us that we must attend to a “range of attitudes,” to “that complicated web of attitudes and feelings which form an essential part of the moral life as we know it”? Why

doesn't he talk about attending to the importance that we attribute to certain other people's attitudes and intentions?

To see what is going on here, we should consider a distinction that Strawson draws. Strawson distinguishes between "the attitude (or range of attitudes) of involvement or participation in a human relationship," on the one hand, and "the objective attitude (or range of attitudes)," on the other (Strawson 1962, 52). This is a distinction between two orientations that we can take towards someone, two different stances. "To adopt the objective attitude to another human being," Strawson writes, "is to see him as, perhaps, an object of social policy; as a subject for what, in a wide range of sense, might be called treatment; as something certainly to be taken account, perhaps precautionary account, of; to be managed or handled or cured or trained; perhaps simply to be avoided..." (Strawson 1962, 52). This objective stance is "profoundly" opposed to its participant counterpart; it precludes "the range of reactive attitudes which belong to involvement or participation with others in inter-personal human relationships" (52).

I suspect that Strawson, in our central passage, has in mind this distinction. Indeed, he even mentions the objective stance there. He tells us that we must attend "to that complicated web of attitudes and feelings which form an essential part of the moral life as we know it, and *which are quite opposed to objectivity of attitude*" (my emphasis). Moreover, he goes on to use the phrase "range of attitudes," which earlier appears when he introduces the "attitude (or *range of attitudes*) of involvement or participation in a human relationship" (my emphasis).

Unfortunately, Strawson doesn't offer much of a positive characterization of the participant stance. It's clear, though, that he at least takes this stance to leave us prone to respond to other people with the range of reactive attitudes. It's also clear that Strawson takes this stance to be tied up with our participation in relationships, and he suggests that we abandon it when we come to see someone as "warped or deranged or compulsive in behavior," as "incapacitated in some or all respects for ordinary inter-personal relationships" (see Strawson 1962, 52-55). I propose, then, that we understand the participant stance as being constituted by the basic concern that Strawson first draws our attention to. That is, to take the participant stance toward someone is to be basically concerned about that person's quality of will, to be invested in their attitudes and intentions such that one is prone to respond to that person's behavior with the reactive attitudes.

Where does all of this leave us? I'm suggesting that Strawson, in "Freedom and Resentment," is engaged in connective analysis and that he takes our basic concern, or the participant stance, to be a core structural feature of our responsibility practices. In this respect, as I understand Strawson, the particular reactive attitudes to which we're prone aren't themselves as important to understanding moral responsibility as the basic concern underlying them. Or to put it more directly, what it means to be morally responsible, on my reading, is ultimately a matter of the basic concern that our responsibility practices express. The nature of this basic concern provides a certain shape to our responsibility practices and to the concepts they involve; it makes certain sorts of social and psychological facts about other people and their actions relevant for how we engage and

how we ought to engage with them, and this gives sense to our notion of moral responsibility.

5. The Concern-Based Construal of Strawson's Reversal¹⁴

In previous work, I've referred to my reading as the "concern-based construal" of Strawson's reversal (Beglin 2018). Before considering how this construal of the reversal helps shed light on Strawson's reconciliatory project in "Freedom and Resentment," let me say a bit more about it.

First, to better see the idea, it might help to consider a simpler example than our responsibility practices. Imagine a particular restaurant's kitchen practices. Let's say that these practices express a concern with efficiently producing food for customers. If these practices are an expression of this concern, then two things follow. First, the kitchen's practices will reflect the concern. The kitchen's practices, that is, will be, on the whole, organized around efficiently producing food for customers, even if they do so poorly. Second, critical reflection on these practices or of decisions made by participants in them will be based on this concern. Thus, if the practices are organized in a way that fails to efficiently produce food, they're criticizable, and if a participant makes a decision contrary to this concern, that person is criticizable. What it means to be "good" or "bad" in the context of this practice will be a matter of the concern with efficiently producing food.

For Strawson, similarly, our responsibility practices are an expression of a certain concern, and they thus reflect this concern and are answerable to it. Crucially, though, the

¹⁴ The material for this section is drawn from Beglin 2018.

concern that our responsibility practices ultimately express, for Strawson, is “basic.” This marks an important difference between the concern expressed by our responsibility practices and the concern in the foregoing kitchen example. The latter concern is best understood in terms of other concerns about, e.g., running a profitable business. These further concerns give reason to organize the restaurant’s kitchen around efficiently producing food for customers, and they might also give reason to organize it differently if circumstances change. The basic concern expressed by our responsibility practices, on the other hand, is simply an essential part of being the kind of social creatures we are.

This point about the *basic* concern is crucial to understanding the Strawsonian reversal. This concern isn’t merely deeply held or important in some community. Rather, it is a concern that forms the general structure for how we (humans) think about moral responsibility; it is the concern that ultimately explains why we have responsibility practices. In this respect, it is like the “human commitment to inductive belief-formation.” It’s “natural, original, non-rational, in no way something we choose or could give up” (Strawson 1962, 64, footnote 7). Thus, for Strawson, while different communities might hold each other responsible for different things in different ways, these differences aren’t as significant for understanding the concept of moral responsibility as what we share. For we share, according to Strawson, a common evaluative standpoint when it comes to responsibility, which makes the same sorts of

general considerations salient for us. Our variable responsibility practices express the same basic concern, and so they're subject to the same general standards.^{15,16}

The idea, then, is that we can only recover what it means to be morally responsible by attending to the basic concern expressed by our responsibility practices, the concern that constitutes the evaluative standpoint (the “participant stance”) from which we hold people responsible. This marks one way in which Strawson thinks that both the optimist and pessimist go wrong. Both, he thinks, “seek, in different ways, to overintellectualize the facts” (Strawson 1962, 64). Specifically, both seek to justify our responsibility practices as a whole. This, though, is to attempt to understand responsibility from outside of the evaluative standpoint that gives sense to it in the first place. As Gary Watson puts Strawson’s point, “the participant standpoint marks the boundaries within which critical assessment of the correctness of our practices makes

¹⁵ Compare this to Strawson’s response to the worry that his discussion of the reactive attitudes might reflect “local and temporary features of our own culture”: “an awareness of variety of forms should not prevent us from acknowledging also that in the absence of *any* forms of these attitudes it is doubtful whether we should have anything that *we* could find intelligible as a system of human relationships, as human society” (Strawson 1962, 65).

¹⁶ This also helps explain why Strawson privileges our basic concern over non-basic—culturally cultivated—social concerns. Our responsibility practices might reflect certain culturally relative concerns, but our basic concern shapes the general structure of those practices. They provide, as Watson (2014) puts it, a “normative framework,” within which our responsibility practices, including the non-basic concerns they might involve, must operate. In this respect, our non-basic, culturally relative concerns are answerable to our basic concern in a way our basic concern isn’t answerable to our non-basic concerns.

sense. Rejections of this perspective on the whole leave us with no relevant ground on which to stand” (Watson 2014, 22).¹⁷

6. The Concern-Based Construal and Strawson’s Reconciliatory Project

To accomplish his aim of reconciling the optimist and pessimist, Strawson must go beyond the foregoing criticism of them. In this section, then, I’d like to consider how he attempts to reconcile these two disputants. To reconcile the two, recall, Strawson needs to provide a foundation for our responsibility practices that allows him to capture the sense of desert that the pessimist feels is missing from the optimist’s account of responsibility, while showing that this notion doesn’t entail freedom of the sort that is incompatible with determinism. The concern-based construal of Strawson’s reversal sheds light on how this reconciliatory project is supposed to work. It also reveals the limits of what Strawson suggests in “Freedom and Resentment,” a fact that I’ll turn to in the next section.

Let’s begin by considering in more detail what the pessimist feels is missing from the optimist’s picture of responsibility. This will give us a better sense of the burden on Strawson.

The optimist takes our responsibility practices to be based on their efficacy “in regulating behaviour in socially desirable ways” (Strawson 1962, 46). On such a view, conditions like mental illness, coercion, and ignorance are relevant to determining whether someone ought to be held responsible because they’re relevant to determining

¹⁷ This is what Watson calls Strawson’s “normative framework argument.” See Watson 2014, 21-24 for helpful discussion. I take my concern-based construal of Strawson’s reversal to complement much of Watson’s reading of “Freedom and Resentment.”

whether holding that person responsible will serve any regulatory purpose. The pessimist finds this suggestion alienating. For the pessimist, that is, these conditions are relevant to determining whether someone should be held responsible because they're relevant to determining whether that person *deserves* to be held responsible. And desert, the pessimist insists, entails a sort of "positive" freedom that the optimist's picture leaves out: "a genuinely free identification of the will with the act" (Strawson 1962, 47).

For the pessimist, the optimist's view renders something that should be essential to moral responsibility—this positive notion of freedom, the identification of one's will with one's act—at best incidental to it. Strawson thinks the pessimist is right to worry about this aspect of the optimist's view. The problem, as Strawson sees it, is that the optimist's picture is dominated by an "objectivity of attitude" (see Strawson 1962, 62). For the optimist, that is, those whom we hold responsible are understood as objects of social policy, as something to be "managed or handled or cured or trained." But this isn't how we relate to people when we hold them responsible. In this respect, Strawson believes that the pessimist is attuned to the real basis of our responsibility practices: the basic concern and the participant stance that it constitutes. The problem with the pessimist, though, according to Strawson, is that she refuses to stop there. Strawson thus hopes to show that we can understand desert and the relevant sort of positive freedom that it involves in terms of our basic concern and the participant stance. We needn't go further than our social sentimental nature, for Strawson.¹⁸

¹⁸ In fact, Strawson's point is stronger than this. He argues that it is *useless* to try to seek a justification from outside of our social sentimental nature. This is tied up with the foregoing argument, which I alluded to at the end of the last section. Strawson's

To show how we can recover the sense of desert that the pessimist feels is missing from the optimist's account, Strawson focuses particularly on the negative reactive attitudes. He writes:

Indignation, disapprobation, like resentment, tend to inhibit or at least limit our goodwill towards the object of these attitudes, tend to promote an at least partial and temporary withdrawal of goodwill; and they do so in proportion as they are strong; and their strength is in general proportioned to what is felt to be the magnitude of the injury and to the degree to which the agent's will is identified with, or indifferent to, it. (Strawson 1962, 63)

Here, Strawson seems to be putting us in a position to understand two elements of desert. The first element is a retributive connotation that the notion of desert often carries with it. Oftentimes, that is, talk of desert is attended by the idea that it is good for one to suffer, or, at least, by the idea that we would be justified in allowing one to suffer. This element of desert, Strawson seems to be suggesting, comes out of the fact that the negative reactive attitudes involve a withdrawal of goodwill. Strawson thus goes on to explain that “the preparedness to acquiesce in that infliction of suffering on the offender which is an essential part of punishment is all of a piece with this whole range of attitudes towards the offender which are in question here” (63).

This retributive element, though, isn't essential to the pessimist's worry, and it doesn't even seem essential to the notions of moral responsibility and desert. We can sensibly talk, for example, of someone deserving forgiveness or gratitude or admiration, and when we do we evidently make no reference to its being good for that person to

argument here is complicated, though. See Strawson 1985 for a more detailed discussion of his “naturalistic” approach. For some different ways of understanding what Strawson is getting at, see Russell 1992, Coates 2017, and Campbell 2017. Here, I'll focus on the reconciliatory part of “Freedom and Resentment.”

suffer. It's important not to be misled, then, by the particular reactive attitudes that Strawson focuses on. Throughout "Freedom and Resentment," Strawson focuses on the negative reactive attitudes. His larger point, though, doesn't seem to rely on this emphasis.

More important for Strawson's reconciliatory project is the second element of desert, the sort of freedom that it entails. Strawson elucidates this aspect of desert by pointing out that our resentment and indignation are attuned to the degree to which one's will is identified with, or indifferent to, what happened. This helps to make room for the sort of positive freedom that concerns the pessimist. What seems to matter is the degree to which the injury reflects one's attitudes or intentions. In other words, it seems that we can identify one's act with one's will insofar as that act reflects one's practical rationality or one's volition.¹⁹ It's important to emphasize, though, again, that the point needn't be made only in terms of resentment and indignation. Strawson's larger point is that the whole range of reactive attitudes are attuned to this sort of positive freedom, and this is because these attitudes grow out of a more general concern that we have about the quality of will reflected in people's behavior.

¹⁹ This way of putting the point is admittedly vague. For our purposes, I doubt we need to make it more precise. I suspect Strawson merely intended to gesture at something intuitive: that there is a real sense in which we can talk about someone's actions reflecting their agential perspective. Interestingly, identificationist views of moral responsibility have become a popular strategy for compatibilism since Strawson's paper was published. See, for instance, Frankfurt 1971 and Watson 1975 for important early statements of this type of view. Strawson's view might be seen as a precursor to these views. For an interesting discussion of the relation between Strawson and Frankfurt's work, see McKenna 2005.

Strawson's discussion thus establishes a connection between his social sentimental basis for our responsibility practices and the sort of positive freedom that the pessimist worries is missing from the optimist's picture. But is this sort of freedom compatible with determinism? It seems so. As Strawson writes, early in "Freedom and Resentment," "people often decide to do things, really intend to do what they do, know just what they're doing in doing it: the reasons they think they have for doing what they do, often really are their reasons and not their rationalizations" (Strawson 1962, 47). And Strawson suggests that this sort of freedom is not ruled out by the thesis of determinism: "For it would not follow from that thesis that nobody decides to do anything; that nobody ever does anything intentionally; that it is false that people sometimes know perfectly well what they are doing" (47).

This positive sort of freedom thus represents one condition on responsibility—one's action must, to some degree, reflect one's agency. Strawson bolsters this idea by pointing to a variety of considerations that lead us to excuse people from responsibility. These are considerations, Strawson explains, that

give occasion for the employment of such expressions as 'He didn't mean to', 'He hadn't realized', 'He didn't know'; and also all those which might give occasion for the use of the phrase 'He couldn't help it', when this is supported by such phrases as 'He was pushed', 'He had to do it', 'It was the only way', 'They left him no alternative', etc. (Strawson 1962, 50)

Such considerations are relevant to responsibility, for Strawson, because they indicate circumstances under which what happened doesn't reflect one's agency. If someone pushes me into you, for instance, my knocking you down doesn't reflect my practical rationality or volition: I didn't want to knock you down; I made no decision about it; etc.

Thus, when I say, “I couldn’t help it, I was pushed,” I’m making clear that my agential capacities weren’t exercised in bringing about what happened, and so what happened doesn’t reflect any attitude or intention on my part. Similarly, if I’m coerced into doing something and I explain, “I couldn’t help it, they left me no alternative,” I’m trying to express that whatever I did doesn’t reflect what I really wanted and that my reasons for doing it weren’t offensive reasons—I wasn’t being careless or cruel. For Strawson, though, determinism is not among the considerations that bear on the question of whether what happened reflects one’s agency. That is, even if the world is determined, it still seems possible that one can act with ill will or indifference; one’s ill will or indifference would simply be determined (see Strawson 1962, 53).²⁰

Still, as Strawson recognizes, whether someone is morally responsible isn’t *only* a matter of the quality of will reflected in that person’s action; it isn’t only a matter of whether what happened reflects one’s agency. There remains the question of whether the person is a morally responsible agent in the first place. To better understand this level of assessing responsibility, Strawson considers pleas such as ‘He has been under very great strain recently’, ‘He was acting under post-hypnotic suggestion’, ‘He’s only a child’,

²⁰ Some theorists have raised worries about Strawson’s invocation of the idea that whether someone “could help it” is relevant to responsibility. Surely, the thought goes, if determinism obtains, then no one can help what they do (see, e.g., Fischer 2014 and Kane 2016). We cannot understand the relevance of phrases like “couldn’t help it” in the abstract, though, and, for Strawson, what we mean by such phrases, the sorts of considerations we have in mind with them, is a matter of the basic concern. See McKenna 2005 for a helpful discussion of Strawson’s “argument from excuses” and the need to understand this argument in terms of Strawson’s focus on quality of will. Interestingly, McKenna connects Strawson’s line of argument to Frankfurt’s (1969) idea that what matters for moral responsibility is what the agent *actually* does. In this respect, Strawson’s ideas might complement an actual-sequence approach to responsibility (cf. Fischer and Ravizza 1998).

‘He’s a hopeless schizophrenic’, and ‘His mind has been systematically perverted.’ Such pleas, Strawson suggests, “invite us to view the agent in a different light from the light in which we should normally view one who has acted as he has acted (Strawson 1962, 51). Particularly, such pleas invite us to adopt the objective stance towards that agent. And as we saw earlier, for Strawson, these pleas invite us to adopt the objective stance because they suggest that the agent is “*incapacitated* in some or all respects for ordinary interpersonal relationships” (55; Strawson’s emphasis).

Why, though, is this incapacity for ordinary relationships relevant to whether we should abandon the participant stance in favor of its objective counterpart? Strawson isn’t entirely clear about this; he doesn’t explicitly address this question. Still, it seems that we can reconstruct an answer based on what he says about the basic concern: if someone is incapable of participating with us in ordinary relationships, then we cannot stand in those relationships with that person; and it thus makes sense that we shouldn’t have the basic concern about that person’s attitudes and intentions, since this concern is somehow tied up with our participation with others in such relationships.

There are reasons to be unsatisfied with Strawson’s account of morally responsible agency, reasons that I’ll turn to in a moment, and which bear on his reconciliatory project. But first, given his account, why doesn’t Strawson think determinism is relevant to morally responsible agency? His argument begins by noting a feature of determinism: it is a global thesis; if determinism obtains, then *everyone* is determined. And according to Strawson, “it cannot be a consequence of any thesis which is not itself self-contradictory that abnormality is the universal condition” (54). In other

words, if the thesis of determinism is true, the fact that we're determined will be a feature of all human relationships; it will be ordinary. Determinism, then, can't impair one's capacity to participate in ordinary human relationships, and so this further condition on responsibility doesn't implicate determinism.²¹

There are good reasons to worry about Strawson's views concerning morally responsible agency. Let's turn, then, to some of Strawson's shortcomings. Ultimately, I hope to show that my concern-based understanding of Strawson's project helps to illuminate where Strawson goes wrong, and thus helps to point the way towards a more illuminating Strawsonian account of morally responsible agency.

7. Strawson's Lacuna

Of the various aspects of "Freedom and Resentment," perhaps none have raised as many questions for theorists sympathetic to Strawson as his account of morally responsible agency and its requirements. I think that theorists are right to worry; this is an aspect of Strawson's thought that requires far more development.

Paul Russell (1992), for example, raises worries for Strawson's idea that morally responsible agency is a matter of a statistical sense of "ordinariness" or "normality." This idea is reflected in Strawson's argument that determinism cannot be relevant to responsible agency. Strawson writes:

“...the participant attitude, and the personal reactive attitudes in general, tend to give place, and, it is judged by the civilized, should give place, to objective attitudes, just in so far as the agent is seen as excluded from ordinary adult human relationships by deep-rooted psychological abnormality—or simply by

²¹ See Pamela Hieronymi's unpublished manuscript, "Freedom, Resentment, and the Metaphysics of Morals," for a very helpful, in-depth discussion of this aspect of Strawson's argument.

being a child. But it cannot be the consequence of any thesis which is not itself self-contradictory that abnormality is the universal condition.” (Strawson 1962, 53-54)

The potential self-contradiction here seems to be the idea that if determinism undermines responsibility, then this is because it renders *everyone* abnormal. But everyone being abnormal is only a contradiction if we understand abnormality to be a statistical notion, determined by what most people are like. Is morally responsible agency really a matter of such a statistical sense of abnormality? Russell argues it isn't. We can imagine, he explains, a world in which most people, even all people, become incapacitated, losing whatever capacities are normally associated with morally responsible agency, perhaps because of the spread of some disease. There seems to be nothing incoherent or self-contradictory about this possibility. And in such a case, it might very well be that the statistically “abnormal” person is a morally responsible agent. Russell takes this to suggest that what is relevant to morally responsible agency is a notion of *capacity*, not statistical ordinariness or normality (Russell 1992, 298-299).

Of course, as we've seen, Strawson himself sometimes talks in terms of capacity. We exempt people from moral responsibility, he writes, when they are “*incapacitated* in some or all respects for ordinary interpersonal relationships” (Strawson 1962, 55; his emphasis). But Russell seems right that the relevant capacities aren't tied to ordinary relationships, at least in a statistical sense. If we drop the statistical notion of ordinariness, though, Strawson doesn't have much of a proposal.

We could try to understand “ordinary” here in a non-statistical sense. That is, we could take Strawson to be drawing our attention to a particular type of recognizable

interpersonal relationship, one that implicates our responsibility practices. Indeed, Strawson at one point suggests that our proneness to the reactive attitudes is a result of our being “involved in interpersonal relationships as we normally understand them” (Strawson 1962, 54). The idea, then, might be that if someone is incapable of participating with us in the relevant kind of interpersonal relationships, relationships as “we normally understand them,” then that person isn’t appropriately regarded as a morally responsible agent, and we should take the objective stance towards them. This seems to fit with Strawson’s idea that the basic concern and the participant stance are connected to our participation in relationships with one another. The problem, though, is that Strawson doesn’t illuminate how we “normally understand” such relationships, or what they involve. Strawson gives us examples of relationships that are characterized by an objectivity of attitude, like the parent-child relationship and the patient-therapist relationship; however, he doesn’t provide a basis for distinguishing these interpersonal relationships from those that are characterized by the participant stance, by a proneness to the reactive attitudes, by responsibility. And because Strawson doesn’t elucidate the relevant kind of interpersonal relationships he has in mind, it is difficult to understand what makes the difference between morally responsible agents and those agents whom we ought to exclude from such relationships (see Bennett 1980, 34-36; Wallace 1994, 29).

Is there any room to develop Strawson’s ideas here? Some theorists have worried that there isn’t. R. Jay Wallace writes, “we do not have an independent concept of an interpersonal relationship suitable to play the required role in the account” (Wallace

1994, 29). Jonathan Bennett, too, despairs, admitting that he “cannot find any [way of determining the relevant kind of relationship] that will do the job” (Bennett 1980, 36). I think that there is a way forward here, though. It rests on better understanding of the relation between the basic concern and the sort of interpersonal relationships Strawson draws our attention to.

In this chapter, I’ve been suggesting that we understand Strawson as taking moral responsibility to be based on a basic concern that we have about people’s attitudes and intentions. This concern constitutes Strawson’s participant stance. And Strawson, we’ve seen, suggests that this concern is tied up with our participation in interpersonal relationships. Strawson, though, isn’t entirely clear about the relation between this concern and those relationships. At points, it seems as though we have the basic concern *because* we stand in the relevant kind of relationships. But there is another way we can understand the relation between this concern and those relationships. It could be that the basic concern is more fundamental. Perhaps, that is, we should understand the relevant type of relationships in terms of this concern. If this is the case, then we can use the basic concern to understand both those relationships and the kind of agency they require. The basic concern, in other words, can illuminate a certain type of relationship at the core of social and moral life, the type of relationship that implicates our responsibility practices, and it can thus illuminate the sorts of capacities required for participation in this sort of relationship, the capacities required for morally responsible agency.

We have, then, at least a way of locating the kind of relationships that are relevant to our responsibility practices and morally responsible agency. The idea is to characterize

them in terms of the basic concern. But this requires a better understanding of the basic concern. And this is where we can locate the central lacuna in Strawson's account. Strawson understands the basic concern only in terms of its object, in terms of what it is a concern *about*. The problem is that we can care about the same object in different ways. Thus, a parent might not regard their child as responsible, but they might still care about the attitudes and intentions reflected in their child's actions. We can say the same thing about a therapist. In this respect, Strawson doesn't provide a full account of the basic concern. Specifically, Strawson doesn't tell us *how* we care about people's attitudes and intentions when we have the basic concern about them. And unless we understand this particular *mode* of concern, we cannot understand morally responsible agency, and we cannot distinguish the kind of relationships that implicate our responsibility practices from other sorts of relationships, which are characterized by an objectivity of attitude, an absence of the basic concern.

What I think the Strawsonian project needs, then, is a fuller account of the basic concern, the concern that constitutes the stance that we take towards morally responsible agents, and which characterizes the kind of relationships responsible agents enter into with each other. To understand morally responsible agency, we need to account for both the basic concern's object *and* the nature or mode of that concern.

8. Conclusion

In this chapter, I've developed a concern-based construal of Strawson's reversal thesis, according to which what it means to be morally responsible is determined by the nature of the basic concern expressed by our responsibility practices. This concern-based

understanding of the reversal sheds light on Strawson's reconciliatory project, and particularly on the strategy that he takes in "Freedom and Resentment." It also sheds light on why Strawson fails to establish an illuminating account of morally responsible agency.

Drawing on Strawson, then, I believe that we can only understand morally responsible agency in terms of the basic concern. This represents a move away from the predominant Strawsonian approach, which understands morally responsible agency in terms of the reactive attitudes themselves. In the next chapter, I'll argue that my concern-based Strawsonian approach is more promising than this attitude-based alternative. I'll then, in chapter three, develop an account of the nature of the basic concern. This will fill in what is missing in Strawson's project. Finally, in chapter four, I'll consider some of the possible implications that my account of the basic concern has for our understanding of morally responsible agency.

It's worth noting that this dissertation is ultimately aimed at better understanding morally responsible agency, and so I won't focus on the debate about whether moral responsibility is compatible with determinism. While much of this chapter had to do with Strawson's reconciliatory project, then, my discussion of this aspect of Strawson's project was largely in the service of unpacking his more general strategy for understanding moral responsibility. In what follows, then, I will shift focus away from Strawson's reconciliatory project, and instead will discuss morally responsible agency on its own terms.

Chapter 2

Two Strawsonian Strategies

1. Introduction

In the last chapter, I developed a particular reading of Strawson's so-called reversal thesis. On my concern-based construal of this thesis, what it means to be a morally responsible agent depends on the basic concern expressed by our responsibility practices. From this thesis follows a strategy for accounting for morally responsible agency: account for it in terms of the nature of this basic concern. In this chapter, I'll argue that this concern-based Strawsonian strategy is more promising than the predominant approach that Strawsonian theorists have taken when accounting for morally responsible agency.

The predominant Strawsonian approach has been developed in terms of the reactive attitudes that characterize our responsibility practices—attitudes like resentment, indignation, and gratitude. On this attitude-based Strawsonian strategy, one accounts for morally responsible agency by accounting for the nature of these attitudes, how they engage the person at whom they're directed.¹ While there are certain virtues to this approach, I worry that the reactive attitudes are the wrong level at which to base a theory

¹ It's worth noting that many theorists understand morally responsible agency in terms of the reactive attitudes insofar as they subscribe to the idea that a morally responsible agent is an apt or appropriate object of such attitudes. While this is no doubt a Strawsonian idea, it's worth distinguishing theorists who subscribe to this idea from theorists who pursue a Strawsonian *approach* to morally responsible agency. As I understand it, the Strawsonian approach is characterized by the idea that the relevant sense of aptness or propriety depends in some way on the nature of our practices of holding responsible. For a similar point, see Watson 2014, 15-16. See also Menges 2017, for a distinction between *weak* and *strong* Strawsonianism.

of morally responsible agency. Rather than account for morally responsible agency in terms of the reactive attitudes, I'll argue, we'd do better to focus on something more fundamental, namely the concern underlying them, the concern that leaves us susceptible to the reactive attitudes in the first place—the basic concern. By accounting for morally responsible agency in terms of this concern we can capture what is compelling about the attitude-based Strawsonian strategy, while avoiding its problems. Further, the concern-based Strawsonian strategy puts us in a better position to understand the distinctive social and moral significance of being and being regarded as a morally responsible agent.

2. The Attitude-Based Strawsonian Strategy

Of the ideas associated with P.F. Strawson's "Freedom and Resentment," perhaps none have been as widely taken up as his notion of "reactive attitudes." Many theorists, even those who distance themselves from other aspects of the Strawsonian program, have followed Strawson in characterizing our responsibility practices in terms of these attitudes, in terms of resentment, indignation, gratitude, guilt, and so on. Still, as we saw in the last chapter, Strawson introduced these attitudes in order to make a larger point about responsibility. "Only by attending to this range of attitudes," he writes, "can we recover from the facts as we know them a sense of what we mean, i.e. of *all* we mean, when, speaking the language of morals, we speak of desert, responsibility, guilt, condemnation, and justice" (Strawson 1962, 64; his emphasis). I've already discussed my preferred way of understanding this passage. Most Strawsonian theorists, however, have developed Strawson's idea differently. While these theorists have offered a variety of distinctive views about morally responsible agency, they've developed those views

through a common strategy, accounting for morally responsible agency by attending to the *nature* of the reactive attitudes.

Gary Watson was perhaps the first person to develop the Strawsonian approach in this way. In his influential (1987) paper, “Responsibility and the Limits of Evil,” Watson argues that Strawson’s account of responsibility, from “Freedom and Resentment,” is incomplete. Specifically, Watson argues, Strawson fails to account for the conditions that morally responsible agency requires. Strawson, according to Watson, tells us what sorts of agents we don’t consider morally responsible—young children, for instance, or severely mentally ill people—but he doesn’t tell us *why* we exempt such agents from responsibility. Strawson doesn’t tell us “what kind of explanations exempt” or “how this works” (Watson 1987, 263).²

To fill this lacuna in Strawson’s account, Watson suggests that we understand the reactive attitudes as “forms of communication, which make sense only on the assumption that the other can comprehend the message” (Watson 1987, 265). Building on Strawson’s suggestion that the reactive attitudes express a demand for goodwill or regard, that is, Watson proposes that the reactive attitudes aim at communicating a moral demand for reasonable regard. And such demanding, Watson explains, “presumes understanding on the part of the object of the demand” (264). Intelligibly communicating a demand to another person, in other words, requires that the other person be capable of understanding what is being demanded of them. If someone is incapable of comprehending the demand

² As we saw in the last chapter, Watson is a part of a long line of theorists who are sympathetic with Strawson but disappointed by his remarks about morally responsible agency. See also Bennett 1980; Russell 1992, 298-301; and Wallace 1994, 27-30.

for reasonable regard that the reactive attitudes communicate, then, Watson holds, that person is an unintelligible object of that demand and, thus, an unintelligible object of the reactive attitudes. In such a case, the reactive attitudes “lose their point as forms of moral address” (265).

Many theorists have followed Watson, adopting some variation of this communication proposal. There are interesting differences between these theorists. What they share, however, is more important here. All agree that the reactive attitudes are forms of moral address, which communicate some kind of moral message, and they all argue that this puts certain constraints on what is required for someone to be a morally responsible agent.³ Watson and his followers thus all employ what I’m calling the “attitude-based Strawsonian strategy,” accounting for morally responsible agency in terms of the nature of the reactive attitudes.

R. Jay Wallace, in his influential book, *Responsibility and the Moral Sentiments*, also employs this strategy. Wallace’s account begins with the following thesis: “S is morally responsible (for action x) if and only if it would be appropriate to hold s morally responsible (for action x)” (Wallace 1994, 91). Wallace suggests that we construe the sense of “appropriate” in this thesis as concerning “fairness.” Thus, for Wallace, one is a morally responsible agent if and only if one has the capacities that would make it fair—

³ See, for instance, Darwall 2006; Shoemaker 2007, 2011, 2015; Watson 2011; and Macnamara 2015a, 2015b. Relatedly, some theorists, particularly Matthew Talbert (2012) and Angela Smith (2013), have developed an alternative to Watson’s communication proposal, which borrows from Watson’s basic idea. These theorists take blame, or resentment, to also, or primarily, protest morally offensive claims that are implicit in wrongdoing, and they argue that this feature of blame puts constraints on morally responsible agency.

reasonable—for one to be held morally responsible. And for Wallace, these are the capacities that make it fair, or reasonable, to respond to one with the reactive attitudes.⁴ What capacities are these? According to Wallace, the reactive attitudes, by which he means only resentment, indignation, and guilt, are essentially ways of holding people to moral obligations. And so, he argues, it is only fair, or reasonable, to respond to someone with the reactive attitudes if that person is capable of grasping and applying the moral reasons that underlie the obligations those attitudes hold that person to (see Wallace 1994, chapter 6). Wallace, like Watson, ultimately accounts for morally responsible agency by appealing to the nature of the reactive attitudes.

Finally, other theorists, like Michael McKenna (2012) and David Shoemaker (2015, 2017), have accounted for morally responsible agency in terms of the “fittingness” of reactive attitudes. This approach begins with an insight about emotions. Emotions are commonly thought to have a cognitive component, representing the world in some way, and they are said to be unfitting when they misrepresent the world. Fear, for example, is often taken to be a response to danger. If this is the case, the thought goes, when one fears something, one represents that thing as dangerous, and, so, if one fears something that isn’t in fact dangerous, one’s fear isn’t fitting; it misrepresents the world. Similarly, the foregoing theorists argue, the reactive attitudes are emotional responses to something, and they are thus only fitting when they are directed at whatever that is. To what, then, do the reactive attitudes respond? The predominant suggestion is “quality of will,” understood in terms of another’s regard. Thus, McKenna (2012), for instance, holds that

⁴ Wallace unpacks these points in his “methodological interlude” (see Wallace 1994, 84-109).

the reactive attitudes respond to someone's action expressing regard or disregard, and so he concludes that morally responsible agency requires the capacity for one's actions to express such regard or disregard.⁵

Strawsonian theorists have thus developed a number of different accounts of morally responsible agency. They arrive at these different accounts through a common strategy, though. All account for morally responsible agency in terms of what the reactive attitudes are, taking the nature of the reactive attitudes to explain what capacities are required for morally responsible agency. For all of these theorists, that is, whether someone is a morally responsible agent depends on whether that person can be an appropriate object of the reactive attitudes, and whether this person can be an appropriate object of the reactive attitudes is a matter of some feature that those attitudes have. Thus, Watson and those who follow him argue that the reactive attitudes communicate a moral message of some kind, and they therefore hold that morally responsible agency requires that one can be an intelligible object of such communication; Wallace argues that the reactive attitudes hold people to moral obligations, and so takes it that morally

⁵ Shoemaker (2015) roughly agrees with McKenna, but he complicates the picture by distinguishing three types of responsibility: accountability, which concerns regard and disregard; attributability, which concerns the quality of someone's character; and answerability, which concerns the quality of someone's rational judgments. I'm not sure that we should distinguish types of responsibility like this. Here, though, I don't intend to weigh in on this issue. I suspect that Shoemaker's "tripartite" view might fare better than its rival attitude-based theories of responsibility; however, I don't think that it avoids the issues that I'll raise for the attitude-based strategy in the next section. If we should distinguish different types of responsibility, I believe we'd do better to distinguish them in terms of different ways of being concerned about other people's actions and attitudes, rather than in terms of the different reactive attitudes those types of responsibility involve (I make a suggestion like this in Beglin 2018, note 15). For more on the different "faces" of responsibility, see Watson 1996; Fischer and Tognazzini 2011; and Shoemaker 2011. For critical discussion, see Smith 2012, 2015, and Watson 2018.

responsible agency requires one be able to grasp and apply the reasons underlying those obligations; and fittingness theorists, like McKenna, argue that the reactive attitudes construe someone as having acted with a particular quality of regard, and so hold that morally responsible agency requires one be capable of acting with that quality of regard.

All of these theorists, then, employ the attitude-based Strawsonian strategy, accounting for morally responsible agency in terms of the nature of the reactive attitudes that characterize our responsibility practices. Why adopt this attitude-based approach? There are, I think, two clear reasons. First, the attitude-based Strawsonian strategy is naturalistic. It provides, that is, a way of accounting for the conditions on being morally responsible that appeals only to facts about human beings and the sorts of attitudes with which we're prone to respond to each other. The attitude-based strategy thus demystifies moral responsibility in precisely the sort of way that Strawson advocated in "Freedom and Resentment." Second, this strategy has a certain intuitive appeal. If we understand holding people responsible in terms of the reactive attitudes, then it seems natural to think that the propriety of holding someone responsible is a matter of the propriety of those attitudes. It seems natural to think, in other words, that we can understand what is required for morally responsible agency by reflecting upon the attitudes that go into treating someone *like* a morally responsible agent.

3. Worries for the Attitude-Based Strawsonian Strategy

While there has been plenty of discussion of the particular Strawsonian theories canvassed above, the strategy that those theories employ—the attitude-based Strawsonian strategy—hasn't itself received much critical attention. Where there has been critical

discussion, it has largely concerned the more general Strawsonian commitments underlying the strategy. Some theorists, for instance, have questioned whether the reactive attitudes are essential to our responsibility practices; others the Strawsonian strategy of deriving an account of moral responsibility from the nature of those practices.⁶ These lines of critique, of course, have implications for the attitude-based Strawsonian strategy, but they ultimately concern underlying features of any Strawsonian approach. Here, then, I'd like to discuss the attitude-based Strawsonian strategy more directly. I'm sympathetic with the general Strawsonian approach to moral responsibility. I doubt, though, that the attitude-based Strawsonian strategy is the best way to develop it.

As a starting point, it is worth emphasizing that advocates of the attitude-based Strawsonian strategy generally attempt to account for morally responsible agency in terms of some feature that the reactive attitudes share, as a class. This task is complicated, though, by the great variety of possible reactive attitudes. Consider, for instance, the diversity of reactive attitudes to which Strawson adverts, in "Freedom and Resentment." Strawson writes of resentment, indignation, guilt, gratitude, moral approbation, anger, shame, forgiveness, hurt feelings, "the sort of love which two adults can sometimes be said to feel reciprocally," and even "such phenomena as feeling bound or obliged" (Strawson 1962, 48, 52, 57). Other theorists, influenced by Strawson, have suggested that we extend the class of reactive attitudes to include attitudes that Strawson doesn't mention, such as contempt (Mason 2003) and trust (Helm 2014).

⁶ For recent discussions of this first line of critique, see Sher 2006, Scanlon 2008, and Wallace 2011. For recent discussions of the second line, see Todd 2016, Shoemaker 2017, and Beglin 2018.

If we adopt such inclusive understandings of the reactive attitudes, then this puts certain pressures on the attitude-based Strawsonian strategy, at least if that strategy involves finding a feature that the reactive attitudes all share. In short, the wider the variety of possible reactive attitudes, the more difficult it is to specify a feature common to all of them, on which to base one's account of moral responsibility. Take, for example, Watson's communication view. It is far from clear that all of the reactive attitudes that Strawson mentions are communicative, addressing a demand for some kind of response from their object. Forgiveness, for instance, isn't obviously communicative in this way, and neither are hurt feelings, guilt, or the feeling of obligation. Still, even if all of the reactive attitudes are communicative, it remains unlikely that attitudes like love, forgiveness, and hurt feelings call for or demand the same kind of response as, say, resentment or gratitude. And this creates the possibility that some of these attitudes, which either aren't communicative or which communicate distinctive messages, implicate different agential capacities than others do.

Something similar can be said about the fittingness of these attitudes. Are all of the reactive attitudes that Strawson mentions even emotions, akin to fear? Forgiveness, reciprocal love, and the feeling of obligation don't obviously fit this paradigm. And the reactive attitudes that do seem to fit this paradigm nevertheless aren't obviously emotional responses to the same thing. What, for instance, distinguishes anger from resentment? I'm not sure there is consensus here. However, some theorists take resentment to be moralized in a way that anger isn't. Wallace, for example, suggests that resentment responds to someone's violating a moral obligation, whereas anger responds

to someone's merely violating some (perhaps non-moral) norm (Wallace 1994, 39). If this is right, then these reactive attitudes might implicate different agential capacities in their object. Likewise, compare hurt feelings to resentment. Both might be responses to how someone regards us, or perhaps to how someone regards something important to us; but this isn't a particularly illuminating level of description. More commonly, as we've seen, resentment is taken to respond to someone's action reflecting *disregard* for us (see, e.g., Shoemaker 2011, 2015; McKenna 2012). This, though, makes resentment seem different from hurt feelings. Hurt feelings don't seem to respond to disregard *per se*. One's feelings might be hurt by perfectly respectful rejection. And again, this leaves open the possibility that resentment implicates different agential capacities than hurt feelings.

To be clear, I don't mean to defend any particular understanding of the foregoing attitudes. My point is only that adopting a wide class of reactive attitudes, like Strawson does, puts certain pressures on the attitude-based Strawsonian strategy, at least if that strategy involves finding a common feature among the reactive attitudes. These pressures haven't gone unrecognized by Strawsonian theorists. Wallace, for instance, in developing his account of moral responsibility, explicitly addresses them. An "inclusive interpretation of the reactive attitudes," Wallace writes, "frustrates any attempt to provide an informative account of what unifies this set of emotions as a class" (Wallace 1994, 11). In light of this problem, Wallace adopts a narrower interpretation, according to which only resentment, indignation, and guilt are reactive attitudes.⁷ These three attitudes, Wallace argues, are unified because they all essentially hold people to

⁷ See Wallace 1994, 25-33 and Wallace 2014, 121-122 for discussion.

obligations. Wallace then goes on to account for morally responsible agency in terms of this feature of his restricted class of reactive attitudes.

While not every attitude-based theorist has been as explicit as Wallace with respect to the foregoing pressures, it seems they've nevertheless recognized the difficulties such a wide class of reactive attitudes poses. Indeed, most Strawsonian theorists operate, at least implicitly, with a narrow interpretation of the reactive attitudes. Watson, for example, develops his communication proposal in terms of the "negative" reactive attitudes.⁸ Likewise, McKenna explicitly follows Wallace, prioritizing resentment, indignation (or moral disapprobation), and guilt (McKenna 2012, 64-66). Even Shoemaker (2015), who admirably makes a point of incorporating positive reactive attitudes into his account, and who distinguishes three types of responsibility, still operates with a narrower class than Strawson's, focusing only on three pairs of attitudes: anger and gratitude, disdain and admiration, and regret and pride.⁹

One way for proponents of the attitude-based Strawsonian strategy to handle the foregoing difficulty, then, is to narrow the class of reactive attitudes. This makes it easier

⁸ "The *negative* reactive attitudes," Watson writes, "express a moral demand, a demand for reasonable regard." He continues: "Since the *negative* reactive attitudes involve this demand, they are not (as fully) appropriately directed to those who do not fully grasp the terms of this demand" (Watson 1987, 264; my emphases).

⁹ Distinguishing different types of responsibility, like Shoemaker, alleviates some of the pressure that is put on the attitude-based strategy by the great variety of possible reactive attitudes. Ultimately, I'm not sure that we should distinguish different types of responsibility like this, but the criticisms I'm developing in this section should apply to the attitude-based strategy in either case. If we distinguish different types of responsibility, the criticisms can be tuned to those particulate types of responsibility, rather than to responsibility writ large. (See also note 5, above.) It's worth adding that it would be notable if the attitude-based Strawsonian strategy must distinguish different types of responsibility to remain plausible.

to specify a feature that the reactive attitudes all share, and this in turn makes it easier to develop an attitude-based account of moral responsibility.

I worry, though, that narrowing the class of reactive attitudes like this creates a new problem for the attitude-based Strawsonian strategy. Earlier, I mentioned that part of what is so intuitive and appealing about this strategy is the thought that we can understand what is required to be a morally responsible agent by reflecting on the nature of the attitudes that go into treating someone *like* a morally responsible agent. In narrowing the class of reactive attitudes, however, attitude-based theorists risk divorcing their account of morally responsible agency, and indeed their understanding of the reactive attitudes, from other important attitudes that are involved in regarding and treating someone as responsible. Why privilege resentment over gratitude? Indignation over moral approbation? Why neglect attitudes like hurt feelings and forgiveness? The feeling of obligation? Insofar as a morally responsible agent is one whom it is appropriate to treat as responsible for her conduct or attitudes, insofar as such an agent is a fitting participant in our responsibility practices, it seems worrisome to account for morally responsible agency in terms of only some of the ways we treat people as responsible, in terms of only some of the attitudes that characterize our responsibility practices. Relatedly, we might worry that the decision to focus on some of these attitudes and not others reflects a prejudgment about responsibility, a prejudgment that doesn't seem to accord with the spirit of the Strawsonian project, which understands responsibility in terms of these attitudes, not the other way around.¹⁰

¹⁰ Thanks to Agnieszka Jaworska for helping me to appreciate this latter point.

So far, I've been assuming that the attitude-based Strawsonian strategy must proceed by identifying a feature that the reactive attitudes share, as a class. And indeed, most attitude-based theorists seem to operate under this assumption. Still, we might wonder about whether the attitude-based strategy must proceed in this way. Why not account for morally responsible agency in terms of all of the essential features that the reactive attitudes have, regardless of whether they share those features with other reactive attitudes? Thus, for instance, gratitude might have feature A, and resentment might have feature B, and hurt feelings might have feature C—and perhaps features A, B, and C all provide requirements for morally responsible agency. Adopting such an aggregative approach would allow one to avoid the pressures put on the attitude-based strategy by the great variety of possible reactive attitudes.

I doubt, though, that we should take this aggregative attitude-based approach. It seems to leave us without a unified understanding of morally responsible agency. For instance, if the reactive attitudes have different features that implicate different capacities, then it is possible that an agent has some of these capacities but not others. And in this case, it is possible that an agent can be, in principle, an appropriate object of resentment, say, but not gratitude. This, however, begins to fracture our notion of morally responsible agency. Would such an agent be a morally responsible agent? In either case, the answer isn't satisfying. If the agent isn't a morally responsible agent, then this seems odd. After all, the agent can be, in principle, an appropriate object of resentment. If the agent *is* a morally responsible agent, though, it seems merely to be because this agent can sometimes be an appropriate object of *some* reactive attitude. But is this really why an

agent counts as a morally responsible agent? This doesn't seem to be a very illuminating way to think about morally responsible agency; there doesn't seem to be any real cohesion here, only a sort of checklist of reactive attitudes and capacities.

Even in the best-case scenario the aggregative variant of the attitude-based approach seems worrisome. It could be, that is, that the individual features of the reactive attitudes implicate the same set of capacities. Thus, any agent who can be, in principle, an appropriate object of resentment might also happen to be capable of being an appropriate object of the other reactive attitudes, and any agent who cannot be an appropriate object of resentment might also happen to be incapable of being an appropriate object of the other reactive attitudes. Further, this correlation, between being an appropriate object of a particular reactive attitude and being an appropriate object of the other reactive attitudes, might hold for each reactive attitude. But in this case, we might wonder: whence the harmony? Why does a heterogeneous group of attitudes, with distinctive features, implicate the same capacities? This seems to be more than mere coincidence; it suggests that something deeper is going on. But the aggregative attitude-based approach doesn't appear to be in a position to explain whatever that is.

I doubt, then, that we should adopt the aggregative variant of the attitude-based Strawsonian strategy. But the problems that it faces point us towards a larger worry for the attitude-based strategy itself. This larger worry arises in virtue of the fact that the various reactive attitudes can themselves have numerous features. Even if we can identify a shared feature among the reactive attitudes, then, it isn't clear that the attitude-based Strawsonian strategy thereby avoids the problems facing its aggregative variant.

Say, for instance, that the reactive attitudes are all communicative, morally addressing their object with the same moral message or demand. Might at least some of those attitudes nevertheless have other essential features, features that aren't shared among all of the reactive attitudes? It seems so. After all, even if the reactive attitudes belong to a singular class, they are still different attitudes with different characters; they must be distinguished somehow. If the reactive attitudes have distinctive essential features, though, do these distinctive features have implications for morally responsible agency?

Perhaps they don't. Perhaps only the features that are shared among all of the reactive attitudes are relevant for understanding morally responsible agency. But I don't see why. A number of theorists, for example, have suggested that the blaming reactive attitudes, like resentment and indignation, protest—stand up against, challenge—morally offensive claims that are implicit in wrongdoing.¹¹ Say that this is true, but add that all of the reactive attitudes, blaming and non-blaming alike, share a communicative feature. In this case, it isn't obvious that the communicative feature, which is shared among all of the reactive attitudes, is more relevant to understanding morally responsible agency than the blaming reactive attitudes' protest feature. If what drives the attitude-based Strawsonian strategy is the idea that we can understand morally responsible agency by thinking about the way we engage morally responsible agents through the attitudes we have towards them, then it seems like all of the ways that we engage morally responsible agents through those attitudes should be relevant.

¹¹ See, e.g., Hieronymi 2001, Talbert 2012, and Smith 2013.

But if all of the features of the reactive attitudes are relevant to morally responsible agency, even those features that aren't shared among the class of reactive attitudes as a whole, then the attitude-based Strawsonian strategy faces the same problems as the foregoing aggregative approach. Either the reactive attitudes' different features implicate different agential capacities or they don't. If they implicate different capacities, then, as we saw above, our account of morally responsible agency begins to look problematically fractured and unsatisfying. If they implicate the same capacities, however, then this is curious. Why would a heterogeneous group of attitudes, with different features, which engage their objects in different ways, implicate the same agential capacities? This calls out for an explanation that the attitude-based Strawsonian strategy doesn't seem able to provide.

4. The Concern-Based Strawsonian Strategy: An Alternative

The problem with the attitude-based Strawsonian strategy, I suspect, is that the reactive attitudes are a relatively superficial aspect of our responsibility practices. This isn't to deny their significance for those practices—or, for that matter, for our lives, relationships, and communities; but I suspect that these attitudes and their features are the wrong level at which to base a theory of morally responsible agency. The great variety of attitudes involved in our responsibility practices, along with the great variety of features they might have, is testament to this. We'd do better, I believe, to focus on something more fundamental and less variable than the reactive attitudes themselves. And the concern-based Strawsonian strategy does just that. Rather than focus on the reactive

attitudes themselves, we should focus instead on the concern that leaves us susceptible to those attitudes in the first place.

Before turning to the relative advantages of the concern-based Strawsonian strategy, let me, in this section, say a bit more about it and its relation to the attitude-based strategy.

To begin, it's important to stress that emotions and feelings don't arise out of nowhere. They come about because we're invested in our world in particular ways. I might feel worried for my friend, or disappointed by a choice she makes. I might feel sad to hear something fell through for her, or relieved to hear she got a job. I might feel happiness, frustration, or fear on her behalf. Such emotional turns come about because my friend matters to me; I care about her. If this person were someone I was indifferent about, I wouldn't be prone to these kinds of feelings or emotions with respect to her.¹² And the same is true of the reactive attitudes. They, like all emotions and feelings, are rooted in some kind of investment in our world, what I'm calling, following Gary Watson, the basic concern.¹³

¹² Here, I'm influenced by Bennett Helm's work on caring (see Helm 2001; see also Jaworska 2007). It's worth noting, though, that in this paper I'm employing the terms "care" and "concern" in a broader sense than Helm. For Helm, caring is always oriented towards something's well-being. I mean to use the term to capture, more generally, a sort of emotional stance of investment in something. So, for example, caring about someone, in Helm's sense, is importantly distinct from being attached to someone (see Wonderly 2016). For my purposes, by contrast, attachment can be said to be a form of caring or concern.

¹³ See Watson 2014. There, Watson suggests that we have a basic concern about how people regard one another and a basic demand that to be treated with regard and good will. As we've seen, Watson understands morally responsible agency in terms of the basic demand: "to be a responsible agent is to be someone whom it makes sense to subject to such a demand" (Watson 2014, 17). My thought, *pace* Watson, is that

This distinction between the emotions and feelings to which we're prone, on the one hand, and the concern that they're embedded in, on the other, is important to keep in mind, because it suggests a further distinction between two levels at which we can assess people's emotions and feelings. The first level focuses on the particular emotions or feelings that one might feel in some circumstance. Thus, one might criticize my fear of flying by explaining that flying isn't actually dangerous, citing various statistics about how people are more likely to die on their drive to the airport than in a plane crash; or one might criticize my resentment as pointless when its object isn't capable of responding to or understanding what it communicates. This first level of assessment contrasts with a second. The second level focuses on the concern that leaves one prone to certain emotions and feelings in the first place. For example, I might feel uneasy about my niece and nephew roughhousing near a vase in my living room. I might squirm when I see them begin to play, or I might audibly gasp. And my sister might criticize this reaction: "Oh, calm down. If they break that thing, they'd be doing you a favor. It's butt ugly." Here, what is in dispute isn't whether my worry or fear is in fact tracking a danger posed to the vase; what is in dispute is whether I should be concerned about the vase in the first place. Of course, my concern about the vase might have nothing to do with its beauty. Perhaps it is of sentimental value. But what's relevant here is that if I were to explain this to my sister, I'd be explaining why I'm concerned about the vase, not why I perceive a threat to it. I'd be justifying my feelings by justifying why the vase is an object worthy of the concern underlying them. In this respect, I wouldn't only be justifying the fears that my

responsible agency is better understood in terms of the basic concern, which, Watson himself notes, seems to underlie the basic demand in any case.

niece and nephew are eliciting at this particular moment; I'd also be justifying other kinds of emotions or feelings that I might have about the vase, such as the sense of relief that might attend my niece and nephew moving on to a new game.

The concern-based Strawsonian strategy locates questions of morally responsible agency at this second level of assessment. The concern-based Strawsonian strategy, that is, understands morally responsible agency at the level of the concern that leaves us prone to the reactive attitudes, and not at the level of those attitudes themselves. The idea is that our responsibility practices, and particularly the reactive attitudes undergirding them, are embedded in a particular form of concern, a particular way of being invested in what other people do or in the attitudes that they have. Whether someone is appropriately treated as a participant in our responsibility practices, then, whether someone is, in principle, an appropriate object of the reactive attitudes that characterize those practices, is a matter of whether it is appropriate to have this basic concern towards them. And so to account for morally responsible agency, on the concern-based strategy, we must first account for the nature of this concern, this way of being invested in people's agency.

5. The Concern-Based Strategy Avoids the Attitude-Based Strategy's Problems

We're now in a position to see why the concern-based Strawsonian strategy avoids the problems faced by the attitude-based strategy. Those problems revolved around the great variety of possible reactive attitudes, along with the seeming variability between them. While the attitude-based Strawsonian strategy struggles in light of such variety and variability, the concern-based strategy can easily countenance it. Concerns, after all, can manifest in different ways. As we saw above, my concern for my friend

might manifest in happiness, frustration, disappointment, sadness, relief, worry—a whole gamut of emotions and feelings, which revolve around my friend’s significance to me. Similarly, our basic concern can leave us prone to a variety of reactive attitudes. These reactive attitudes represent different ways of engaging with the significance that we ascribe to someone’s actions and attitudes when we have the basic concern toward them. This suggests that the reactive attitudes *should* be multifarious. Still, insofar as these attitudes reflect the same concern, they’ll reflect the nature of that concern. In this respect, the concern-based approach puts us in a position to put the different features that the reactive attitudes might have into a larger context, providing a means for showing how these features, though distinctive, are nevertheless connected.

At this point, it might help to have an account of the basic concern on the table. In the next chapter, I’ll develop such an account. Let me briefly preview that account here, though, for the purposes of illuminating these points about the concern-based strategy. To be clear, I don’t take the concern-based strategy to depend on my particular understanding of the basic concern. I also don’t mean to fully develop my account of the basic concern just yet. What follows is a rough adumbration.

Strawson takes the basic concern to be a concern about people’s attitudes and intentions.¹⁴ As we’ve already seen, though, identifying the object of a concern isn’t

¹⁴ It’s worth noting that we can spell out the object of this concern in different ways. For instance, Strawson sometimes sounds as though he has in mind a concern about people’s attitudes and intentions *towards us*, but at other times he seems to have in mind a more general concern about people’s attitudes and intentions. Theorists also, as we’ve seen, spell out these ideas in terms of how people regard us or other people. In the next chapter, I’ll spell out, in more detail, how I understand the object of this concern. For the purposes of the example, though, I’ll remain vague.

sufficient for accounting for that concern, because we can care about the same things in different ways. My concern for my friend and her well-being, for instance, might very well encompass her attitudes and intentions. Her attitudes and intentions, that is, might matter to me because they bear on how she is doing. Such a concern would surely leave me susceptible to certain feelings and emotions; however, it is doubtful that it itself would leave me susceptible to the sorts of feelings and emotions characteristic of our responsibility practices. Likewise, a psychiatrist might care about her patient's attitudes and intentions, including his good or ill will, in a way that doesn't leave her in anyway susceptible to resentment or indignation. The psychiatrist's concern might merely be clinical; she might just want to diagnose her patient or help her patient cope with things.^{15,16}

I take the main question facing the concern-based Strawsonian strategy, then, to be about the way in which we care about people's attitudes and intentions when we have

¹⁵ There is a distinction to be drawn here between being concerned about one's attitudes and intentions themselves and being concerned about *knowing* what one's attitudes and intentions are. The basic concern appears to be a concern of the first type, as does one's concern about one's friend's attitudes and intentions. Arguably, the psychiatrist's concern might be an epistemic concern; her interest might just be in knowing what her patient's attitudes and intentions are, so that she can diagnose her patient and treat her patient. Thanks to Agnieszka Jaworska for drawing my attention to this subtlety.

¹⁶ The psychiatrist example, and the general point I'm making here, draws on Jonathan Bennett's worries about Strawson's view. As we've seen, Strawson links our proneness to the reactive attitudes to our participation with people in relationships. Bennett argues, though, that we can share some relationships, like psychiatrist-patient relationships, that are importantly interpersonal but that don't leave us prone to the reactive attitudes. The problem, as Bennett puts it, is that "the terms 'interpersonal relation', 'participation', etc. do not, unaided, mark off the territory which Strawson wants to delimit" (Bennett 1980, 35; see also Wallace 1994, 2014 and Watson 1987). Putting the point in Bennett's terms, then, we might say that the notions of "concern," "importance," "mattering," also do not, unaided, mark off the territory which Strawson wants to delimit.

the basic concern toward them. In what way are morally responsible agents' attitudes and intentions important to us? In what way do they matter? Strawson, as we saw in the last chapter, isn't clear on this point. But more generally, this represents a major lacuna in our theorizing about the moral psychology of responsibility. While many theorists have taken the reactive attitudes to revolve around people's attitudes and intentions, few have characterized this in terms of a concern about those attitudes and intentions. As a result, I believe, theorists have typically focused on the *object* of the basic concern—people's attitudes and intentions—and not on the way those attitudes and intentions matter to us when we have this concern, when we're prone to respond to someone with the reactive attitudes.

As I understand it, when we have the basic concern about someone's attitudes and intentions, we take those attitudes and intentions to bear in a certain way on our own evaluative outlook, on our own view about what sorts of attitudes and intentions that we should have or that we should perceive as appropriate in the relevant context. When we have the basic concern towards someone, in other words, we ascribe a kind of practical authority to their attitudes and intentions; we take their attitudes and intentions to be capable of making a claim on us, to be capable of challenging or affirming our perspective about what's appropriate. In this respect, the basic concern is a concern about someone's attitudes and intentions as the attitudes and intentions of a peer, or a fellow participant, in some social context.

The idea is that our responsibility practices, and particularly the reactive attitudes that characterize them, revolve around this concern. If this is right, then the reactive

attitudes will reflect this way of being concerned about people's attitudes and intentions. Consider, for example, the blaming reactive attitudes. As we've seen, some theorists take the blaming reactive attitudes to protest offensive claims that are implicit in wrongdoing, while others take those attitudes to communicate a moral message or demand. Both features of these attitudes can be understood in terms of the foregoing account of the basic concern. Begin with protest. When someone treats us with indifference or disregard, their attitude might challenge our perspective about our own moral importance, about the kinds of attitudes that are appropriate to have towards us. And this seems to be because we ascribe a kind of practical authority to the other person's attitudes and intentions. Our resentment, then, can be understood as standing up against the claim that the other person's attitudes make on us, the implicit suggestion that we ought to accept the propriety of the other person's indifference or disregard (see Hieronymi 2001). Likewise, a number of theorists have suggested that the blaming reactive attitudes aim at getting the other person to recognize the wrongness of what they've done (see, e.g., Macnamara 2015a and Shoemaker 2015, 103-112). Such a communicative aim makes sense if our resentment is undergirded by the basic concern, as I've suggested construing it. If we have such a concern, then it will matter to us that the other person recognize the wrongness of what they've done because that person's attitudes carry the weight, the authority, of a peer's. There is thus some social urgency to their rebuking their own wrongdoing.¹⁷

¹⁷ There is a deep connection here, then, between communicative and regulative views of the blaming reactive attitudes. See McGeer 2010, 2014, and 2015.

Say, then, that the blaming reactive attitudes have both a protest and a communicative feature. We've already seen that such a plurality of features poses problems for the attitude-based Strawsonian strategy, particularly since these features aren't obviously shared with other reactive attitudes. The problem here is that the attitude-based strategy attempts to account for morally responsible agency in terms of the reactive attitudes' features themselves; it takes morally responsible agency to be explained by those features. The concern-based Strawsonian strategy, on the other hand, can sidestep such issues precisely because it focuses on the more fundamental phenomenon: the concern that our reactive attitudes, including their features, presuppose. Thus, while hurt feelings may not protest or communicate, at least not in the foregoing way, they nevertheless might be understood as being undergirded by the basic concern, as I've construed it. They certainly seem to be the result of our being vulnerable and open to other people's attitudes about us or about the relationships we share with them. Likewise, many theorists use the protest feature of resentment to make sense of forgiveness, which doesn't itself protest, but which can be understood as a response to a past action's ceasing to make a claim on us that merits protest (Hieronymi 2001; Martin 2010).¹⁸

The thought, then, is that rather than account for morally responsible agency in terms of the nature of the reactive attitudes, we should account for morally responsible agency in terms of the basic concern that leaves us susceptible to those attitudes and that informs their nature. On the concern-based Strawsonian strategy, that is, morally

¹⁸ In the next chapter, I'll expand on these remarks about the way my account of the basic concern can illuminate the various reactive attitudes to which Strawson adverts in "Freedom and Resentment."

responsible agency is about the way we care about another person's attitudes and intentions when we hold that person responsible; it is about the significance that we ascribe to them. If my account of the basic concern is correct, then, a morally responsible agent can be understood as a kind of peer, a fellow participant in some social context. Morally responsible agency will thus require one be capable of operating as such a peer. This might mean that one has certain general evaluative and rational capacities, for instance, along with the more substantive capacity to appreciate (or perhaps come to appreciate) the particularities of the social context in which one's responsibility is being considered.

I'll discuss this particular concern-based theory of morally responsible agency over the next two chapters. At this point, I merely hope to use the sketch of this theory to help bring out what is distinctive about the concern-based Strawsonian strategy, particularly in contrast with its attitude-based counterpart. Whereas the attitude-based strategy explains morally responsible agency in terms of the reactive attitudes and the way they engage their objects, the concern-based strategy focuses instead on the concern underlying those attitudes, the concern that leaves us susceptible to and that informs the nature of those attitudes in the first place.¹⁹ And this allows the concern-based Strawsonian strategy to avoid the problems faced by the attitude-based strategy.

¹⁹ The concern-based strategy I've developed here might also be contrasted with Bennett Helm's recent work on responsibility. In *Communities of Respect*, Helm extends his previous work on the nature of caring into the realm of morality. Here, I don't have the space to fully explore the intricacies of his view. It's worth noting, though, that while forms of caring (reverence and respect) play an important role in Helm's thought, he doesn't take responsible agency to be explained in terms of the propriety of such caring itself, as I'm proposing. For Helm, rather, someone is a responsible agent, a member of a

6. Some Other Virtues of the Concern-Based Strawsonian Strategy

The attitude-based Strawsonian strategy thus appears to face certain issues that the concern-based Strawsonian strategy avoids. This, though, isn't the only reason to prefer the concern-based approach. Before concluding, I'd like to proffer two more reasons in its favor. First, the concern-based strategy shares the attitude-based strategy's virtues, and, second, it seems better positioned to capture the distinctive social and moral significance of being and being regarded as responsible.

Earlier, I suggested that the attitude-based Strawsonian strategy had two virtues. The first concerned its naturalism, and the second concerned its intuitive appeal. The concern-based Strawsonian strategy shares both of these virtues.

Consider the attitude-based Strawsonian strategy's naturalism. The attitude-based strategy demystifies responsibility in precisely the sort of way that Strawson advocated in "Freedom and Resentment," providing a way to account for moral responsibility that appeals only to facts about human beings and the sort of attitudes with which we're prone to respond to each other. The concern-based Strawsonian strategy is similarly naturalistic. On the concern-based Strawsonian strategy, one's account of morally responsible agency needs only to appeal to facts about humans and the human concerns that underpin our reactive attitudes. The concern-based Strawsonian strategy thus, like its attitude-based

community of respect, bound by its norms, "just in case other members of that community treat her as a member in their reactive responses" (Helm 2017, 145). I worry that such a view is open to problematic forms of relativism, and I suspect that Helm would do better to put more emphasis on the intelligibility of the concerns that he takes to be involved in treating someone as a member. Interestingly, Helm comes close to something like the concern-based Strawsonian strategy in reply to a worry about the possibility that his view implies, counter-intuitively, that *anything* could be a responsible agent (cf. Helm 2017, 146-149).

counterpart, takes the notion of moral responsibility to simply be a matter of our human, social sentimental nature.²⁰

The attitude-based strategy's second virtue is its intuitive appeal. It seems natural to think that we can understand what is required to be a morally responsible agent by thinking about the attitudes that go into treating someone *like* a morally responsible agent. Of course, as we've seen, the attitudes that go into treating someone like a morally responsible agent might take various forms, and this complicates things. But we can extend this intuitive line of thought to avoid these complications. Rather than merely think about the nature of the attitudes that go into treating someone like a morally responsible agent, rather than merely think about *how* we treat morally responsible agents, we can think about why we're prone to those sorts of attitudes, why we're prone to treat morally responsible agents in *that* way. This latter idea is precisely what motivates the concern-based Strawsonian strategy. In this sense, the concern-based Strawsonian strategy seems rooted in the same sort of intuitive line of thought as its attitude-based counterpart.

By sharing its virtues and avoiding its worries, then, the concern-based Strawsonian strategy seems more promising than the attitude-based strategy. However, there is still another reason to prefer the concern-based approach. Morally responsible agency seems to carry with it a distinctive sort of significance, and the concern-based Strawsonian strategy seems better positioned to capture that significance than its attitude-based counterpart.

²⁰ I owe the phrase "social sentimental nature" to Watson 2014.

It is a familiar idea that holding someone responsible is a form of respect. Strawson, for instance, suggested that our proneness to the reactive attitudes is a way of regarding someone as a “member of the moral community” (Strawson 1962, 63), and this idea has gained much traction, particularly among Strawsonian theorists (see, e.g., Stern 1974; Darwall 2006, 86; and Shoemaker 2007). Herbert Morris made a similar point in his landmark (1968) paper, “Persons and Punishment.” There, Morris dramatized the value that we ascribe to being held responsible by contrasting our familiar “system of punishment” with a “system of treatment,” wherein all wrongdoing is treated like a symptom of mental illness. The latter system, when spelled out, is chilling in its inhumanity. There thus seems to be a distinctive sort of value to being and being regarded as a morally responsible agent. This value is also reflected in our attitudes towards the loss of such agency. Losing the capacities required for morally responsible agency is tragic, and the prospect of such a loss is terrifying. As Angela Smith puts it, “being held responsible is as much a privilege as it is a burden” (Smith 2007, 269).

What accounts for the distinctive significance that we ascribe to being and being regarded as a morally responsible agent? My suggestion is that the concern-based Strawsonian strategy is better positioned to answer this question than its attitude-based counterpart.

On the attitude-based Strawsonian strategy, whether someone is a morally responsible agent depends on whether that person can be, in principle, an appropriate object of the reactive attitudes, particularly of some feature those attitudes have. The attitude-based Strawsonian strategy would thus seem to explain the distinctive

significance of being and being regarded as a morally responsible agent in terms of the reactive attitudes and how they engage their objects. While thinking about how the reactive attitudes engage their objects no doubt hints at the significance of being and being regarded as a morally responsible agent, however, it doesn't seem to put us in a position to fully understand it. For instance, consider the communication view, which is perhaps the attitude-based theory that is best positioned to say something about the significance of being and being regarded as a responsible agent. Say the reactive attitudes communicate some kind of moral message or demand. Such moral address certainly seems significant. But why is it significant? And in what way? Why, in other words, would it be regrettable to lack the capacities required to be an intelligible object of such moral address? And why is communicating a moral demand or message a distinctive form of respect? To answer these questions, it seems like we need to fit the communication feature of the reactive attitudes into a bigger picture. Again, focusing on the reactive attitudes and their features appears to be too superficial.

The concern-based Strawsonian strategy, I think, already focuses on the bigger picture. On the concern-based approach, morally responsible agency is a matter of the basic concern around which our responsibility practices revolve, the concern that gives rise to our reactive attitudes in the first place. This is a deeply social concern about other people's attitudes and intentions. And because this concern underlies our reactive attitudes, it should shed light on the nature of those attitudes and on how they engage their objects. In this respect, the significance of the reactive attitudes' features can be better understood in terms of the basic concern. Thus, for example, consider, once more,

my preferred account of the basic concern. We've already seen how this can illuminate the communicative nature of the blaming reactive attitudes. Those attitudes might aim at communicating the wrongness of what the other person did—they might aim at getting that person to appreciate the import of their actions—and our attitudes might have this communicative aim precisely because we see that person's attitudes as mattering *for us*, as carrying a kind of practical authority. By fitting this communicative feature of resentment and indignation into a larger framework, the basic concern can thus illuminate the value that we ascribe to morally responsible agents when we morally address them in this way. Such moral address is undergirded by our regarding them as our peers.

The concern-based Strawsonian strategy, then, seems positioned to shed more light on the significance of being and being regarded as a morally responsible agent than its attitude-based counterpart. It also seems to speak more directly *to* that significance. Regarding someone as a morally responsible agent, on the concern-based approach, is a way of being invested in their attitudes and intentions, a way of caring about their agential perspective. And being a morally responsible agent means having the capacities that are required to be appropriately regarded in this way. Thus, when someone lacks the capacities required to be a morally responsible agent, it is because that person lacks the capacities required for one to be appropriately concerned about that person's attitudes and intentions—it is because one shouldn't regard that person's attitudes and intentions as mattering in a particular way. It is easy to see why this could be a regrettable state of affairs. Likewise, it is easy, on this approach, to see why regarding someone as a morally

responsible agent is a form of respect. It is, after all, a way of treating that person as someone that matters.

We can also see this point, about capturing the distinctive significance of being and being regarded as morally responsible, by considering the mistake we make when we hold someone responsible who *isn't* a morally responsible agent. The attitude-based Strawsonian strategy explains this mistake in terms of the propriety of some feature of our reactive attitude. When I resent an infant, for instance, on this approach, the problem is that my resentment isn't fitting or fair or intelligible, because the way it engages its object isn't fitting or fair or intelligible. Thus, Watson and communication theorists might hold that my resentment of this infant isn't intelligible, because she isn't an intelligible object of the moral demand or message that my resentment communicates; and Wallace might suggest that my resentment of the infant isn't fair, because she isn't capable of grasping or applying the reasons underlying the obligations my resentment holds her to; and fittingness theorists like McKenna might suggest that she isn't a fitting object of my resentment, because her actions can't express the disregard that my resentment construes her action as expressing.

These explanations, though, don't seem quite right. That is, the problem with my resentment doesn't seem to be with the resentment *per se*; it seems like it has to do with how I'm relating to this infant in the first place. I shouldn't be prone to resent her, but neither should I be prone to have hurt feelings when she isn't interested in me or to have

gratitude toward her for things that she does.²¹ Being so prone means I'm comporting myself toward this infant inappropriately; I'm inappropriately invested in her attitudes and intentions. And this is precisely the sort of explanation the concern-based Strawsonian strategy gives.

The concern-based Strawsonian strategy, then, not only avoids the worries of its attitude-based counterpart; it also provides a naturalistic, intuitively appealing way of accounting for morally responsible agency, one that seems to better illuminate what it *means* to be and to be regarded as a morally responsible agent.

7. Conclusion

The attitude-based Strawsonian strategy has been, by far, the predominant approach that Strawsonian theorists have taken to accounting for morally responsible agency. In this chapter, I hope to have articulated and motivated my concern-based alternative. I don't mean to suggest that I've provided a knockdown argument against the attitude-based strategy, though. Still, I think that the concern-based Strawsonian strategy is more promising. Once we step away from the reactive attitudes and focus on the concern that leaves us susceptible to them, I believe we can obtain a better sense of what it means to be a morally responsible agent. Ironically, stepping away from the reactive attitudes might also give us a better view of the reactive attitudes themselves, allowing us to understand them, as a class, while embracing, even illuminating, their diversity.

²¹ I might have gratitude *about* something she does, but this isn't the same as feeling gratitude *toward* her *for* something (see Walker 1980-81 and McAleer 2012). Only the latter seems to be a reactive attitude.

What we need is a fuller account of the basic concern, one that goes beyond what I adumbrated above. I'll turn to the task of developing such an account in the next chapter. I'll then, in chapter four, apply the concern-based strategy, exploring some possible implications that my understanding of the basic concern has for how we should understand morally responsible agency.

Chapter 3

The Basic Concern

1. Introduction

Over the course of the last two chapters, drawing on P.F. Strawson, I've suggested that our responsibility practices, and particularly the reactive attitudes that characterize them, revolve around a particular way that we care about, that we're invested in, how people exercise their agency. Following Gary Watson (2014), I've called this the "basic concern." As we've seen, the basic concern is an area of our moral psychology that theorists have largely overlooked, or, at least, left underdeveloped. I suspect that this is due to the tendency to focus instead on the reactive attitudes, and I believe it has been a mistake. Focusing on the basic concern will put us in a position not only to better understand morally responsible agency, but also to better understand our responsibility practices more broadly, including the various reactive attitudes that they involve.

In this chapter, then, I turn my attention to the basic concern, putting forth a particular conception of that concern, one that goes beyond what theorists have previously said about it. This will put me in a position to shed some light on the notion of morally responsible agency in the next chapter. It should be added, though, that the basic concern is worth investigating for its own sake, underpinning, as it does, so much of our social and moral lives.

2. The Basic Concern's Object

How, then, should we understand the basic concern? As I understand it, the basic concern has two aspects, and I'll begin, in this section, by addressing the first of them: its

object. To account for the basic concern, that is, we must account for what it is a concern *about*. This was Strawson's emphasis in "Freedom and Resentment," and subsequent theorists who have discussed the basic concern have followed his lead in taking this aspect as their focus.

Strawson introduces the basic concern early in "Freedom and Resentment," in the same section where he introduces the reactive attitudes. "The central commonplace that I want to insist on," he writes, "is the very great importance that we attach to the attitudes and intentions towards us of other human beings, and the great extent to which our personal feelings and reactions depend upon, or involve, our beliefs about these attitudes and intentions" (Strawson 1962, 48). Later, after introducing a few provocative examples ("simplifications"), Strawson redoubles his point: "These simplifications are of use to me only if they help to emphasize how much we actually mind, how much it matters to us, whether the actions of other people—and particularly of *some* other people—reflect attitudes towards us of goodwill, affection, or esteem on the one hand or contempt, indifference, or malevolence on the other" (49).

At the center of our responsibility practices is thus, for Strawson, a concern about the attitudes and intentions reflected in people's actions, a concern about people's "quality of will." This idea is central to the Strawsonian program, and it gets at something that has resonated even with theorists who distance themselves from other aspects of

Strawson's thought.¹ Still, there remains some question about how precisely we should understand Strawson's idea here.

We might worry, for instance, that conceiving of the basic concern as a concern about the attitudes and intentions reflected in people's actions cannot capture how we're apt to respond to certain agents' negligent or forgetful actions. When my friend forgets about my birthday party, there seems at least to be a sense in which her forgetting doesn't reflect any attitude or intention, or even a "willing," on her part. She didn't *intend* or *will* to forget about the party, and it isn't clear that her forgetting necessarily reflects any specific attitude, e.g. indifference or malevolence, towards me. The case seems different, for example, from a case where I'm aware that saying something will hurt someone's feelings but I indifferently or malevolently say it anyways. Nevertheless, we're apt to respond to people's negligent and forgetful actions with reactive attitudes. I might appropriately resent my friend's forgetting about my birthday party, for instance. This suggests that the basic concern might not be about attitudes and intentions *per se*.²

In fact, most Strawsonian theorists have moved away from talking in terms of the attitudes and intentions reflected in someone's actions. The predominant way theorists

¹ T.M. Scanlon, for example, who is skeptical about Strawson's emphasis on the reactive attitudes, has nevertheless developed an influential "quality of will" view of responsibility. See, e.g., Scanlon 1988, 60-68; Scanlon 1998, ch. 6; Scanlon 2008, ch. 4; and Scanlon 2015. Angela Smith, too, has been critical of aspects of Strawson's project (see Smith 2007), but, following Scanlon, endorses a quality of will view (Smith 2005).

² We might hold that the basic concern is about attitudes and intentions, but we might maintain that it isn't only about the presence of certain attitudes and intentions. Rather, we might hold that the basic concern is also about the absence of certain attitudes and intentions. I won't pursue this intriguing thought here; instead, I'll draw on the way other theorists have typically developed the basic concern. Many thanks to Agnieszka Jaworska, however, for drawing my attention to this alternative approach.

have understood the basic concern, rather, is in terms of “regard.” Michael McKenna, for instance, explicitly distancing himself from the thought that the basic concern is ultimately about attitudes, intentions, or wills, explains, “What [Strawson] had in mind when speaking of a person who in acting shows good will toward others, or instead ill will, is better captured in terms of that person’s regard or concern for others” (McKenna 2012, 59). Similarly, Gary Watson takes Strawson to be suggesting that we “care deeply (and ‘for its own sake’) about how people regard one another” (Watson 2014, 17).

This notion of regard gets at something more general than attitudes and intentions themselves. The idea is a familiar one: what we do, including the attitudes and intentions underlying what we do, reveals something about our evaluative or normative outlook on the matters at hand, about what we take to be valuable or important or worthwhile in the context of our action.³ What we do, in other words, says something about how we regard things.⁴ My decision to cancel our plans so I can spend time with a new romantic interest, for example, might reveal that I regard spending time with my new romantic interest as being more important than our friendship; and when I scoff at someone’s complaint, my

³ To be clear, by “evaluative outlook” here I don’t mean to suggest anything too substantive. I don’t mean to suggest, for instance, that regard is a matter of one’s all-things-considered take on what is good. Here, rather, I have a thinner notion of “evaluative outlook” in mind: our actions and attitudes reflect how the situation we’re facing strikes us; they reflect, that is, our immediate take on the importance, value, or worthwhileness of the features of the situation.

⁴ It is worth noting that we could use the word “attitude” to get at something similar to this notion of “regard” (cf. Scanlon 1988, 160-166). Nevertheless, I suspect “regard” is an improvement on “attitude,” because the latter is ambiguous between specific attitudes we might have—contempt, affection, indifference, malice—and the more general, evaluative comportment towards things that “regard” gestures at.

scoffing reflects the fact that I don't regard that person's complaint as having much, if any, merit.

As McKenna points out, this more general notion of regard can help us better account for cases of negligent and forgetful action. This is because how someone regards something or someone can be revealed not only "in her intentions, decisions, choices, or reasons for action," but also in "what she fails to attend to" (McKenna 2012, 59). Thus, to return to our original example, the very fact that my friend forgets about my birthday might reveal something about her regard for me. If she forgets about my birthday because she is busy with work, for example, it might reflect the fact that she regards her work as more important than me, her friend. And I might find this offensive. Given our relationship, she isn't according me the appropriate regard.

We have to be careful here, though. This isn't to suggest that my friend, in forgetting about my birthday, makes any explicit or conscious judgments about my importance. Nor is it to suggest that my friend's forgetfulness reflects her "real" feelings about me. When my friend eventually realizes that she forgot about my party, she might become overwhelmed with guilt, and she might immediately call me and apologize. The point, rather, is that when my friend forgets about the party, this reveals that, at that moment at least, she was more concerned with her work than our friendship; she acted as though her work were more important. This, presumably, is why she may feel so bad when she remembers the party: at that moment, her work shouldn't have been more important to her than our friendship. It is also reflected in how we often respond to such

negligent or forgetful agents. We often complain that such people don't have their priorities in the right place; they fail to give something the appropriate level of regard.

Many theorists have connected something like this notion of regard to moral responsibility. They've used different vocabularies for this purpose. Jeffrie Murphy, for example, suggests that when someone wrongs us their action communicates a "symbolic message"—"I count but you do not,' 'I can use you for my purposes,' or 'I am here up high and you are there down below.'" For Murphy, our resentment is a response to such messages (Murphy and Hampton, 25). Similarly, T.M. Scanlon (2008, 53-56) and Michael McKenna (2012, Ch. 4) have suggested that praise and blame respond to linguistic-like "meaning" that people's actions express. Pamela Hieronymi (2001) and Angela Smith (2013), on the other hand, have taken blame to respond to "claims" that wrongful actions implicitly make. And still other theorists have talked in terms of praise and blame responding to the "judgments" reflected in someone's actions or attitudes (Scanlon 1998; Smith 2005, 2008, 2012; Talbert 2008, 2012, 2014).

All of these theorists, I take it, are driving at the same general idea, an idea reflected in the basic concern: there is a deep connection between our responsibility practices, on the one hand, and our preoccupation with what people's actions "say"—with the evaluative take on things that people's actions reflect—on the other. In this way, the basic concern is part of a much larger tradition in the contemporary literature on moral responsibility.⁵ Here, I'll continue to talk in terms of the basic concern being a concern about the *regard* reflected in people's actions; however, it could also be put in

⁵ This should be no surprise. The foregoing theorists' views are likely all, in one way or another, partly the product of Strawson's thought.

terms of a concern about the meaning, messages, claims, or judgments people's actions reflect.

So far, I've focused on the notion of regard, but there is a complication here. The basic concern isn't a concern about one's regard simpliciter. It is, rather, a concern about how people regard *something*. How, then, should we understand this "something"?

Above, Watson and McKenna suggested that, for Strawson, the basic concern is a concern about how people regard *each other*—other people. This is a promising start. There is no doubt that we often praise and blame others for how they treat other people (including, of course, how they treat us). Nevertheless, I worry that this proposal is too narrow. After all, we're prone to resent people, or to feel approbation toward them, for a great variety of things, not all of which obviously reflects any particular regard for other people. One might resent or feel indignant toward someone for their mistreating a dog, say, or for their destroying a work of art or damaging some ecosystem. On the other hand, one might feel approbation toward someone for their efforts to rescue a stray cat or for their environmental conservation efforts. These examples suggest that the basic concern goes beyond being a concern only about how people regard other people.

Considerations like these have led McKenna to broaden his understanding of the basic concern's object. He suggests that the basic concern is a concern about "the regard or concern one has toward others (or oneself), *and toward the relevance of moral considerations*, as manifested in one's conduct" (McKenna 2012, 59, my emphasis; see also McKenna 2016, 248). This seems to be a move in the right direction. McKenna is

right to think that we ought to broaden the basic concern in light of the foregoing cases, and works of art, animals, and the environment are all plausibly moral considerations.

Still, there is reason to worry about McKenna's proposal. His appeal to moral considerations raises a number of difficult and contentious questions about the scope and definition of morality. What precisely counts as a *moral* consideration? Does the consideration have to *in fact* be a moral one, or is it enough that the consideration appear moral to the person with the basic concern? And if the latter, what is it for something to "appear moral"? The problem here is that the scope of McKenna's proposal is unclear.

It may seem unfair to demand answers to these questions from McKenna. Surely we can't expect him, in addition to giving an account of responsibility, to provide a clear account of the scope or definition of morality. Most theorists aren't in a position to do this, and, at the very least, it is a contentious matter. Moreover, it seems desirable for a theory of moral responsibility to remain neutral about the scope or definition of morality. Nevertheless, it isn't clear to me that McKenna, or at least someone who accepts his proposal, can remain neutral about these issues.

The problem is that, depending on one's conception of morality, McKenna's proposal may not capture all of the types of cases in which we're prone to hold people responsible. People blame others for many things that aren't obviously immoral and that don't necessarily seem tied to a concern about how people regard other people. For instance, people resent and feel indignation toward others for their bad etiquette, for their failure to conform to certain social norms about how to dress or wear one's hair, and for things like burning flags, using crass language, and engaging in certain types of sexual

activity with certain types of people.⁶ Furthermore, as Scanlon points out, it can matter a great deal to us how someone regards some joint aim we share, and such joint aims aren't always clearly moral (Scanlon 2008, 54). A band member might thus feel indignant about her bandmate's starting a side project because she takes her bandmate's action to reflect a lack of regard for their shared endeavor, the band. Alternatively, one might feel approbation toward someone for their impressive commitment to some joint aim, for their going over and above to make some shared project succeed. And we could also imagine a person praising someone for her "decent" behavior—dressing conservatively, say, or waiting until marriage before she moves in with her beau.⁷

I thus doubt that we should broaden the basic concern's object in terms of moral considerations. This isn't to say that such a proposal cannot possibly be correct. It seems to depend on what one has in mind by "moral." But this is a question it seems best to avoid having to answer, at least for the purposes of providing an account of the basic concern. How, then, can we more plainly incorporate the wide range of things for which we're prone, or potentially prone, to hold people responsible?

Strawson, in a provocative passage, reminds his readers about the connection between the basic concern and our participation with each other in a wide range of social contexts. He writes:

⁶ See Margaret Urban Walker's insightful discussion of the variety of "offenses" we resent people for (Walker 2006, 124).

⁷ I don't mean to endorse any of these reactions. My point isn't a normative one, about the kinds of things for which people should praise or blame each other. Rather, my point is a moral psychological one, about how people do sometimes respond to each other. This fits with the object of my analysis. The basic concern is, after all, a moral psychological phenomenon.

We should think of the many different kinds of relationships which we can have with other people—as sharers of a common interest; as members of the same family; as colleagues; as friends; as lovers; as chance parties to an enormous range of transactions and encounters. Then we should think, in each of these connections in turn, and in others, of the kind of importance we attach to the attitudes and intentions towards us of those who stand in these relationships to us, and the kinds of *reactive* attitudes to which we ourselves are prone. In general we demand some degree of goodwill or regard on the part of those who stand in these relationships to us, though the forms we require it to take vary widely in different connections. The range and intensity of our *reactive* attitudes towards goodwill, its absence or its opposite vary no less widely. (Strawson 1962, 49-50)

Here, Strawson is focused on our concern about how people regard *us*, but the spirit of his thought is nevertheless applicable to our larger question. We exist with each other in a wide variety of ways. We share in practices, communities, and relationships. And it is within these social contexts, against the background of these practices, communities, and relationships, that what we do and how we regard things is significant for other people.

Drawing on this idea, I propose that we understand the basic concern as a concern about how people regard the relevance of what I'll call “socially salient considerations.”⁸ Socially salient considerations are considerations that we take to matter within the social contexts we inhabit. Our relationships, communities, and practices, that is, are characterized in terms of distinctive concerns, expectations, vulnerabilities, and dependencies, and it is in light of these concerns, expectations, vulnerabilities, and dependencies that certain considerations become important—socially salient—within the context of the relevant relationship, community, or practice. Given the concerns and

⁸ If one were so inclined, one could perhaps refer to these considerations as “moral,” broadly construed. It is common, however, for theorists to use “moral” in a narrower sense, and so I think it will be easier to distinguish these considerations with another term.

expectations that I take to characterize our friendship, for instance, I may take my welfare or my feelings to matter in certain ways within the context of our friendship. We can say something similar about other considerations that gain their significance for us from the various ways we're connected to other people—as “sharers of a common interest,” as “members of the same family,” as “colleagues” and “lovers,” as “chance parties to an enormous range of transactions and encounters.”

Crucially, what is socially salient within these contexts will be, at least in part, socially determined.⁹ At the most general level, what is socially salient will be fixed by our shared cultural or societal understandings of the relevant relationships, communities, or practices, including the concerns, expectations, vulnerabilities, and dependencies they involve. So, for instance, we as a culture might have certain general understandings about what is socially salient—what sorts of things should matter—to and for friends, or we might have shared understandings of the sorts of things parents should take into account when raising their children. These shared understandings represent common conceptions of the relevant kinds of relationships, communities, and practices. It's worth noting, though, that what is socially salient can also be determined at a more particular level. Thus, for example, the history of one's friendship or romantic relationship, along with the idiosyncratic features of each party to it, might inform or change what matters to the people within the context of that particular friendship or relationship.

The notion of a socially salient consideration is flexible, then—deliberately so. Given the nature of socially salient considerations, it isn't difficult to see why other

⁹ Here, I draw on Strawson's discussion of “social morality” (see Strawson 1961, 5ff.) and Margaret Urban Walker's *Moral Understandings* (2006b), especially chapter 4.

people, or maybe, more accurately, other people's interests or welfare, can be—perhaps most often are—socially salient considerations. It is hard to imagine a mode of co-existence that doesn't somehow implicate the interests of those involved. Moreover, many of our relationships are premised on various forms of dependency, which put our welfare directly in the hands of others. There is no doubt, then, that in most social contexts others people's interests will usually be taken to matter in some way.

But more can be socially salient than just our interests. After all, many of our relationships, communities, and practices revolve around shared commitments, and these shared commitments often lead us to have certain expectations of each other. Thus, above I gave the example of a band member blaming her band mate for starting a side project. It isn't difficult to imagine this band member's indignation reflecting the fact that her band mate should maintain a certain level of regard for their band—in this context, in other words, their band itself is a socially salient consideration. Similarly, as Scanlon points out, some communities and relationships are organized around an attachment to certain cultural traditions (Scanlon 2008, 54). Members of such communities might come to expect fellow members to regard their traditions with a certain degree of reverence. Consider, e.g., the way many American conservatives have become concerned about (and have responded viscerally to) a perceived “war on Christmas,” or what they perceive as an attack on certain “Christian” ideals they take America to be founded on. Crucially, such traditions may or may not be moral, depending on what one has in mind by “moral.”

This point about the way relationships, communities, and practices can be organized around certain shared commitments or aims helps give sense to one way in

which the environment might become a socially salient consideration. Some communities might revolve around a shared conviction that human beings are part of a larger natural order, which they must respect. Of course, it could also be that a community takes the environment to matter only in virtue of the fact that it is something we share—if one harms the environment, one harms other people. In either case, the environment is socially salient.

Most anything can be a socially salient consideration, then, depending on the form of life some group of humans share. I thus take my social salience proposal to add a level of flexibility that is important when accounting for the basic concern. The basic concern, after all, is a moral psychological phenomenon, which is supposed to underlie our general proneness to the reactive attitudes; it is meant to apply to all human beings, across cultures and times. But human communities and relationships seem capable of taking many forms, and it seems plausible that we might be prone to reactive attitudes towards people for an equally various set of behavior. It is crucial that an account of the basic concern reflect these facts.

3. Going Beyond the Basic Concern's Object

So far, I've developed an account of the basic concern's object, what it is a concern *about*. On my view, the basic concern is a concern about how people regard the relevance of socially salient considerations. But as we've seen in previous chapters, a full account of the basic concern must go beyond its object. We can care about the same thing in different ways, and so specifying the basic concern's object doesn't distinguish it from other forms of concern.

Thus, drawing on the account of the basic concern's object that I've developed, there are different ways we can be concerned about how people regard the relevance of socially salient considerations. A psychiatrist, for example, might be concerned about how I regard the relevance of certain socially salient considerations, such as other people's welfare; but this concern may not leave her prone to any reactive attitudes, or to hold me responsible for my regard. Parents, too, are sometimes concerned about how their children regard the relevance of certain socially salient considerations, but, again, at least depending on the child's age, not always in a way that leaves them prone to resent or feel indignant about their child's regard. And when a sociologist studies how some human population regards the relevance of certain socially salient considerations, this sociologist is certainly concerned, in some sense, about her subjects' regard, but not in the same way as when she resents me for talking loudly in the library and disturbing her efforts to write up her results. How can we distinguish the basic concern from these other ways in which we might be concerned about how people regard the relevance of socially salient considerations?

As a general point, when we care about something, that thing matters to us in a certain way; it has a certain type of significance for us. Thus, for example, my concern for you might be different than your family's concern for you, and this is plausibly because you matter to me in a different way than how you matter to your family. As Bennett Helm puts the point, we always care about something under a description (Helm 2010, 97). I may care about you *qua* fellow person, or *qua* friend, or *qua* reader; your daughter might care about you *qua* parent, your mother *qua* (adult) child, and your

husband *qua* spouse. These different descriptions indicate different ways in which you might matter to someone who cares about you.

We can apply this general point about caring to the basic concern. When we're basically concerned about how someone regards the relevance of socially salient considerations, we care about their regard under some description—it matters to us in some particular way. And evidently, the way that someone's regard matters to us when we're basically concerned about it will differ from the way in which that regard matters for the foregoing psychiatrist, parent, and sociologist. To distinguish the basic concern from these other ways of being concerned about how people regard the relevance of socially salient considerations, then, we need to specify the way in which people's regard matters to us when we're basically concerned about it.

Here we have the second aspect to giving an account of the basic concern. Theorists have focused on determining the basic concern's object, and this is surely essential. But we must also determine the way in which that object matters to us when we're basically concerned about it. It is to this task that I now turn.

4. The Basic Concern and the Participant Stance

The basic concern, then, is a concern about how people regard the relevance of socially salient considerations. In what way, though, does how someone regards the relevance of socially salient considerations matter to us when we're basically concerned about it? And how does this differ from the way such regard matters to the foregoing psychiatrist, parent, and sociologist?

These are difficult questions to answer. To start, we might turn to another difficult question, one that has received more attention from theorists and that I discussed toward the end of my first chapter. Strawson famously contrasted two stances that we take towards others: the participant stance and the objective stance. The participant stance is the stance “of involvement or participation in a human relationship” (Strawson 1962, 52). This stance leaves us prone to the reactive attitudes, to hold people responsible. The participant stance is thus constituted by the basic concern; to take the participant stance toward someone is to be basically concerned about how they regard the relevance of socially salient considerations. The objective stance, on the other hand, is “profoundly opposed” to the participant stance. To adopt the objective stance toward another human being, Strawson explains, “is to see him, perhaps, as an object of social policy; as a subject for what, in a wide range of sense, might be called treatment; as something certainly to be taken account, perhaps precautionary account, of; to be managed or handled or cured or trained; perhaps simply to be avoided” (52). When we take this stance toward someone “all our reactive attitudes tend to be profoundly modified” (52).

The distinction between the participant and objective stance is interesting for our purposes because it parallels the distinction between the basic concern, on the one hand, and, on the other, the type of concerns expressed by our psychiatrist, parent, and sociologist. Our psychiatrist, parent, and sociologist can all be said to be concerned about how some others regard the relevance of socially salient considerations, but their concern is clearly filtered through particular ways of viewing their subjects. Thus, the psychiatrist’s concern, we can say, is filtered through a view of her patient as an object

“for what, in a wide range of sense, might be called treatment.” Likewise, the sociologist might see the human population she studies as something like an “object of social policy,” or perhaps just something to be understood. And finally, oversimplifying a bit, we can say that our parent’s concern about her child’s regard is, at least in an important sense (though perhaps not always, depending on the child’s age), under the guise of her child as a being to be “managed” or “trained.” The concerns exhibited by our psychiatrist, parent, and sociologist therefore all seem to align (more or less) with the objective stance, while the basic concern aligns with the participant stance.

Of course, the distinction between the participant stance and the objective stance is notoriously fraught. Strawson, for instance, seems to equate the participant stance with our participation in ordinary interpersonal relationships. Thus, he writes, “being involved in interpersonal relationships as we normally understand them precisely is being exposed to the range of reactive attitudes and feelings” (Strawson 1962, 54). As Jonathan Bennett (1980) and R. Jay Wallace (1994; 2014) point out, though, it isn’t clear in what sense a parent-child relationship and a psychiatrist-patient relationship isn’t an interpersonal relationship, despite their being marked by the objective stance. It also isn’t clear how we “normally” understand the notion of an interpersonal relationship.

Strawson thus overstates things, or at least obscures them, when he equates the participant stance with involvement in interpersonal relationships as we normally understand them. The participant stance can’t merely be a matter of being “involved” in an interpersonal relationship, and it isn’t clear how we “normally understand” interpersonal relationships. Once we appreciate the relation between the participant

stance and the basic concern, however, we can more clearly describe the situation. The participant stance is a particular *way* of being involved in a relationship with someone. Specifically, we can say that the participant stance is a particular way of being *concerned* about how that person regards the relevance of socially salient considerations—considerations, we might add, that get their significance from the various forms of relationship and community we share with that person. We can thus grant Bennett and Wallace that psychiatrists and parents are importantly involved in interpersonal relationships with their patients and their children, respectively. The point is that the way they're involved in these relationships—the way they're engaged with the other person—isn't the same as how we're engaged with others when we take the participant stance towards them. The parent, for example, is certainly engaged with his young child, but he isn't engaged with his young child in a way (fully) marked by the participant stance. The parent is, rather, at least to an important degree, engaged with his child as a being to be managed or trained.

This isn't yet to answer our initial inquiry, of course, but it provides a helpful starting point. How, we can ask, should we characterize the way we're engaged with other people when we're basically concerned about how they regard the relevance of socially salient considerations? Answering this question will help us characterize the way that person's regard matters to us when we're so concerned.

As a first pass, and taking the notion of the “participant” stance seriously, we might say that when we're basically concerned about how someone regards the relevance of socially salient considerations, we're engaged with that person as a *fellow participant*

in some practice, community, or relationship. In this sense, we're concerned about that person's regard *qua* the regard of a fellow participant. Still, the question remains: What is it to engage with someone as a "fellow participant"?

In an early discussion of the participant stance, Lawrence Stern suggests that engaging with someone as a member of the moral community means remaining open, in principle at least, to that person changing our mind about moral matters (Stern 1974, 75).¹⁰ Stern holds that this idea is crucial to understanding what Strawson is getting at with the participant stance. He develops this thought in dialogue with a remark that Strawson makes about how the participant stance differs from the objective stance.

Strawson writes:

If your attitude towards someone is wholly objective, then though you may fight with him, you cannot quarrel with him, and though you may talk to him, even negotiate with him, you cannot reason with him. You can at most pretend to quarrel, or to reason, with him. (Strawson 1962, 52)

For Strawson, then, the participant stance in some way underlies our reasoning or quarrelling with someone, whereas the objective stance, at least when fully taken toward someone, excludes these activities. Stern makes sense of this by arguing that reasoning and quarreling with someone are crucially connected to seeing that person as a member of the moral community. When one quarrels or reasons, Stern argues, one recognizes "a certain equality between oneself and the other." He continues:

¹⁰ As will come out below, in the next section, we have to qualify Stern's claim, as the notion of "membership" here is ambiguous. Specifically, we might distinguish between one's having standing (and thus protection) within a (moral) community and one's perspective on things mattering for what other people should think in that community.

There is, in general, no point in reasoning unless the other person is capable of seeing reason, getting the point. If he can do that, he can also correct *me* if I am mistaken. (Stern 1974, 75)

Stern applies this point to moral reasoning in particular:

Moral argument presupposes that the other person is capable of responding to the argument. He is, at least potentially, a member of the moral community with myself. He shares my principles or can come to share them. *Or else he may win me over with his.* (Stern 1974, 75; my emphasis)

And quarrelling, according to Stern, involves the same sort of openness to the other person's perspective. "When two people quarrel," Stern explains, "they trade emotion, each submitting to the impact of the other's feelings." This contrasts, Stern notes, with "shooting the other down the minute he opens his mouth." Thus, Stern concludes, "What quarreling has in common with reasoning and moral argument is that all three involve giving the other the opportunity to change us" (Stern 1974, 75).

Stern's ideas provide a helpful starting point for understanding what it is to engage with someone as a fellow participant in some social context. The thought seems to be, roughly, that engaging with someone as a fellow participant involves taking their outlook seriously. It means ascribing a kind of practical authority—"a certain equality"—to their regard, taking how they see things to bear in some way on how we should see them. We can apply this idea to the basic concern. When we have the basic concern, we can say, we ascribe a kind of practical authority to how someone regards the relevance of socially salient considerations. We thus take how that person regards the relevance of socially salient considerations to bear in some way on how *we* should regard their relevance.

Still, this idea needs to be refined. What's the precise nature of this practical authority? In what sense do we take the person's regard to bear on how we should regard the relevance of socially salient considerations?

Stern's discussion, at points, makes this practical authority sound merely epistemological. The other person can "correct" us; they can "win us over with their principles." Is the relevant practical authority, then, a kind of epistemic authority? When we take the participant stance towards someone, when we have the basic concern about their regard, do we ascribe some degree of credence to their outlook? If we adopt this model, the basic concern would involve taking how someone regards the relevance of socially salient considerations to perhaps be the *right* way to regard those considerations' relevance.

This, though, doesn't seem like the best way to conceive of the basic concern. Consider, for instance, the conviction with which we often blame people. We're often quite sure that the person whom we resent or feel indignant toward has failed to appropriately regard the relevance of socially salient considerations. In fact, it seems like our blame is often strengthened by just how wrong we think the other person was to regard something in the way that they did. In this respect, it doesn't seem like holding another person responsible depends on, or arises from, ascribing any credence to how they regard the relevance of socially salient considerations.

The basic concern, then, shouldn't be thought of as reflecting some preoccupation with finding the correct way to regard socially salient considerations. We don't have the basic concern about how someone regards the relevance of such considerations because

we take that person to be a reliable judge of such matters. How, then, should we understand the basic concern? I propose that we understand it in terms of the shared social contexts in which it operates.¹¹ That is, we co-exist with people in normative contexts, against the backdrop of relationships, communities, and practices in which we aren't the sole voices and in which what matters isn't entirely up to us. In such contexts, as Cheshire Calhoun (2016) points out, normative problems are settled by achieving agreement, shared understandings of how to go on together, and not necessarily by achieving the "correct" view through justificatory arguments.

This is not to say that reason and argument can play no role in settling the normative problems that arise in our shared social contexts. The point, rather, is that in such a context our reasoning and arguing aren't driven by an interest in knowledge. They're driven instead by an interest in coming to a shared understanding.¹² We can imagine, for example, a man and a woman who are about to wed. The woman wants to keep her last name; the man wants her to take his. The woman might explain that she rejects this practice as antiquated and sexist; she's not his property. And this might move her fiancé for different reasons. One possibility is that he comes to believe, with her, that the practice is sexist and wrong. But he might also be moved by something else. He might, for instance, come to think that if the matter is so important to his fiancée, then it should be left to her to make the decision, even if he doesn't see the problem—"I didn't

¹¹ What's the relationship between our shared social contexts and the basic concern? Do we have the basic concern *because* we co-exist in shared social contexts? I suspect that it is more complicated than this. I don't, then, mean to imply that our shared social contexts give rise to the basic concern. The basic concern, though, is tied up with our sharing social contexts with one another. It might even be essential to sharing such contexts.

¹² Here, I draw on Walker 2006b and the introduction to Calhoun 2016.

know you felt that way,” he might say, “of course you don’t have to take my name.” Alternatively, things might be more complicated. He might explain, for instance, that he agrees with her reasons but that it will mar his parents’ relationship with her if she retains her name. His traditionalist parents won’t understand her decision, and it is important to him that she has a good relationship with them. In this latter case, it’s worth emphasizing, the man’s reasons don’t have to do with the correctness of his fiancée’s views. If they were having a philosophical discussion about whether we, as a society, should abandon the practice of women taking their husbands’ names, their discussion could have ended after she articulated the point about sexism. This man might thus agree that his parents are wrong to care about whether his wife takes his name. But, as Margaret Urban Walker explains, resolutions to problems that arise in our shared social and moral life are often more like the outcome of a negotiation than a proof (Walker 2006b, 76).¹³

What’s at stake for the foregoing couple, then, is the relevance of certain socially salient considerations: the woman’s relationship with the man’s family, the woman’s independence and self-respect, along with certain social and cultural ideas and traditions about marriage that inform their social situation. The couple is thus negotiating the very terms of their relationship. And this sheds some light on the nature of the shared social contexts in which the basic concern operates. These are contexts that are driven by an

¹³ Earlier, we saw that Strawson takes negotiation to be compatible with the objective stance. Is this a problem for what I’m saying? I don’t think it is. I take Strawson to have in mind the kind of bargaining we might do with a child: if you eat your vegetables, then you can have dessert. This is different from the negotiation of values we see in the couple I’ve been discussing.

interest in intersubjective agreement, an interest in arriving at a shared understanding of how to go on together.

Moreover, this interest in intersubjective agreement isn't simply instrumental. The woman in our foregoing example doesn't only seem interested in her fiancé's outlook because it bears on what she can do or on her well-being. We could imagine, for instance, that her fiancé accedes to her demands, perhaps because he has a headache and doesn't want to keep talking about whether she'll take his last name. Perhaps he waves his hands and says, "Fine. Do whatever you want. I can't think about this anymore." In this case, the woman gets the outcome she wants: she can keep her last name. But how she secures this outcome might reasonably make a difference to her. She might find her fiancé's response disagreeable or unsatisfying, even offensive. And this is presumably because they haven't actually arrived at a shared understanding. His concession isn't a recognition of the importance of her integrity and independence, and this itself might matter to the woman. Depending on how her fiancé regards the relevance of her integrity and independence, this can change the nature of their relationship; it can put her relations with him on "an entirely different footing" (Scanlon 1998, 76).

The thought, then, is that the basic concern occurs against the backdrop of social contexts in which intersubjective agreement matters in its own right. Against the backdrop of such social contexts, one's regard takes on a new significance for us. We take how that person regards the relevance of socially salient considerations to bear on how we should regard the relevance of those considerations in that how we should go on together is a matter of our agreement, a matter of arriving at a shared understanding. The

person's regard can thus make a claim on us; it can challenge or affirm our own perspective about how we should regard the relevance of socially salient considerations. This, I propose, is what it means to see that person's regard as the regard of a fellow participant in some social context. The practical authority that we ascribe to that person's outlook is a reflection of the way we take ourselves to share social space with them. How we go on together is partly a matter of arriving at a shared understanding with them, and so their regard can challenge or affirm how we see things.

We can thus see how the basic concern underlies the foregoing couple's discussion about last names. When the man suggests that the woman keeping her last name will damage her relationship with his parents, he is suggesting that her relationship with his parents is a consideration that matters in the context of their marriage. And he is also suggesting that her relationship with his parents might be more important than the sense of independence or integrity that is tied up with her keeping her last name. His fiancée, of course, might find this suggestion offensive. Indeed, she might feel indignant or resentful at the suggestion; their discussion might turn into a quarrel. If her fiancé regards her independence or integrity with such little regard, this matters for the actual shape of their relationship. There is thus a kind of urgency to their disagreement. This urgency is due to the fact that she sees his evaluative outlook as bearing on what should matter in the context of their relationship, and so she sees his evaluative outlook as challenging how she should regard things.

The foregoing couple represents a relatively intimate social context. But, of course, the basic concern also operates in wider-ranging situations. Earlier, for instance, I

mentioned the tendency of certain American conservatives to respond viscerally to what they perceive as a “war on Christmas.” When these citizens hear someone exclaim “Happy holidays!” rather than “Merry Christmas!” they might swell with indignation. Why? It seems these citizens take such “political correctness” to reflect a lack of regard for certain “American” traditions and values. And this challenges their conception of what sorts of considerations should matter for Americans. They feel compelled to reckon with the “politically correct” person’s regard, that is, as the regard of a fellow American, someone whose outlook bears on a shared understanding of how to go on as an American. On the flip side, this offended citizen might feel emboldened or affirmed by a fellow American whose behavior or reactions express a similar disdain for such “PC” culture.

The case of the conservative American helps to distinguish between two ways in which we might take someone’s regard to matter, particularly with respect to considerations that we see as socially salient. On the one hand, conservative Americans who viscerally resent people for expressing “PC” sentiments are no doubt responding to a trend that they see as threatening. That is, part of the threat that conservatives perceive is tied up with the fact that many people are saying “Happy holidays!” rather than “Merry Christmas!” and so conservatives feel like their country is changing in ways that they find objectionable and alienating. It’s against this background that someone’s “PC” behavior might become particularly offensive for a conservative. But crucially, recognizing the impact that a person’s regard has with respect to this emerging trend isn’t yet to be basically concerned about that person’s regard. That is, a conservative might respond to

an expression of this emerging trend in different ways. She might, for instance, feel regret or frustration or fear; she might nostalgically reflect on the past; she might even think about retreating, hiving herself off in a secluded community that is centered on traditional values. But when the conservative American resents or feels indignant about the “PC” person’s saying “Happy holidays!” she isn’t simply responding to an emerging “PC” culture; she’s responding, rather, to *this person’s* evaluative outlook. Against the background of a certain emerging consensus, then, the conservative American might take how the other person regards the relevance of certain considerations to carry a particular importance, but it is nevertheless seeing the other person as a fellow participant in the first place that leaves her prone to hold that person responsible.¹⁴

What I’m suggesting is that the basic concern ascribes a kind of practical authority to how someone regards the relevance of socially salient considerations. This authority is the consequence of viewing that person as a fellow participant in some social context, as someone with whom we seek a shared understanding of how to go on together. How that person sees things can thus challenge or affirm how we see things; their outlook can make a claim on how we perceive what should matter in the context that we share with them.

I believe that this way of understanding the basic concern can help us understand a point that is implicit in Strawson’s discussion of that concern. Strawson appeals to a number of examples, “simplifications,” in an attempt to motivate “how much we actually mind, how much it matters to us, whether the actions of other people... reflect attitudes

¹⁴ Thanks to Coleen Macnamara and Agnieszka Jaworska for discussion of this point.

towards us of goodwill, affection, or esteem on the one hand or contempt, indifference, or malevolence on the other” (Strawson 1962, 49). These examples are telling what they have in common. Strawson writes:

Thus we may, like La Rochefoucauld, put self-love or self-esteem or vanity at the centre of the picture and point out how it may be caressed by the esteem, or wounded by the indifference or contempt, of others. We might speak, in another jargon, of the need for love, and the loss of security which results from its withdrawal; or, in another, of human self-respect and its connexion with the recognition of an individual’s dignity. (Strawson 1962, 48-49)

Strawson’s examples all have to do with the way in which someone’s regard for us can impact our sense of self in some way. This is, for better or worse, a very recognizable phenomenon. When one’s actions reflect a lack of regard for us or for our interests—when, for example, someone whom we like stands us up—this may lower our self-esteem or make us less self-respecting; it may make us feel less secure about ourselves. On the other hand, when someone’s behavior reflects a high level of regard for us, this may bolster our sense of self, making us feel more secure in our standing vis-à-vis that person, raising our self-esteem or emboldening our self-respect. Strawson’s thought seems to be, then, that the basic concern leaves us vulnerable to others in a particular sort of way.

Capturing this sort of vulnerability seems important for an account of the basic concern—and not only because Strawson sees the basic concern as involving such vulnerability. A vulnerability to others seems tied up with what Stern says about our remaining open to others, for example; and many theorists have tied discussions of the reactive attitudes, particularly resentment, to something like the

vulnerability Strawson is drawing our attention to.¹⁵ Can the account of the basic concern that I've given capture this sort of vulnerability?

I believe it can. Let's begin with a familiar point: that our identities are socially constituted. Peter Railton, for instance, explains, "Our identities exist in relational, not absolute space, and except as they are fixed by reference points in others, in society, in culture, or in some larger constellation still, they are not fixed at all" (Railton 1984, 167). Similarly, Jeffrie Murphy writes, "We are, at least in part, social and socialized products—creatures whose sense of self is so much a part of our social setting that the idea of self-respect or self-esteem cannot be totally detached from the concern with how others (some others) regard and treat us..." (Murphy and Hampton 1988, 93). I believe that we should understand the basic concern as complementing these ideas.

It might help at this point to follow Robin Dillon (1997) in distinguishing between intellectual and experiential understanding. Intellectual understanding, according to Dillon, "involves having beliefs which one has reason to accept as true, then coming by inference to have other beliefs which one takes to be true in virtue of their logical relation to warranted beliefs, where the believing, inferring, and assessing need not engage the emotions." Experiential understanding, on the other hand, involves "experiencing something directly and feeling the truth of what is experienced" (Dillon 1997, 239). This distinction can help us better understand the way that others can impact us when we have the basic concern about how they regard the relevance of socially salient considerations.

¹⁵ See, e.g., Murphy and Hampton 1988, Dillon 1997, Hieronymi 2001, and Walker 2006a.

Earlier, I suggested that we shouldn't think of the basic concern as epistemological. We shouldn't think of it as reflecting a preoccupation with getting things right; we don't necessarily ascribe credence to the evaluative outlooks of people towards whom we have it. In this respect, I doubt that we should think of the basic concern as being essentially connected to our intellectual understanding of how we should regard the relevance of socially salient considerations. There's reason to think, though, that the basic concern is essentially connected to our *experiential* understanding of how we should regard the relevance of those considerations. Peter Railton explains that the "self has reference points beyond the ego, and that which affects these reference points may affect the self in an unmediated way" (Railton 1984, 166). Railton's idea is that our "sense of self," our identity "as *experienced*," is fundamentally relational.¹⁶ Thus, if I lose a loved one, this can have a direct impact on how I experience myself, on my sense of myself as a nephew or son, say. This unmediated relationality helps to shed light on the way in which the basic concern leaves one open to how others regard the relevance of socially salient considerations, the way in which how certain others regard the relevance of those considerations can directly bear on one's experiential understanding of how one should regard those considerations. Insofar as one's relationship with the other person is premised on shared understandings, how the other person regards the relevance of socially salient considerations can impact the nature of those relations and, thus, one's very sense of self.

¹⁶ The quotes in this sentence come from Railton 1984, 167, note 38.

Say, for example, that my friend stands me up, despite knowing that I've been feeling down. In this case, my friend's behavior reflects a lack of regard for my interests, and, as I noted above, in the case of the disagreeing couple, there is a kind of urgency to this. If my friend accords my interests such little relevance, then this changes the actual shape of our relationship; it changes the very nature of my relations with her. And this is because our relationship is premised on a shared understanding of how we should go on together. Drawing on Dillon and Railton, we can further explicate the urgency here. It isn't merely normative. That is, it isn't *merely* that my friend's behavior raises a question about how we should regard the relevance of my interests within the context of our friendship. The urgency is also existential, implicating my very sense of self. The challenge that I feel isn't intellectual, then, but experiential. In the context of our friendship, I experience myself and the importance of my interests partly through my friend. In this respect, the challenge that I feel is personal. My friend's regard can directly impact my sense of how I should regard the relevance of socially salient considerations.

Something similar might be going on with our foregoing conservative, who is worried about certain types of "political correctness." She has an idea of what it is to be an American, and how other Americans regard the relevance of certain socially salient considerations can thus challenge her very sense of self. In this respect, her identity as an American is at stake for her. Likewise, other Americans, who share her disdain, can directly affirm her sense of self. In this shared social context, she experiences herself through certain others.

With theorists like Railton, Murphy, Strawson, and Dillon, then, I agree that a kind of vulnerability to others' regard is, as Murphy puts it, "an ineliminable part of the human condition" (Murphy and Hampton 1988, 93).¹⁷ Furthermore, I think that the basic concern is tied up with this vulnerability. To be clear, though, this vulnerability is compatible with one's having strong convictions about what should be regarded as socially salient. Consider our conservative American. Her strong reactions to "PC" culture surely reflect strong beliefs about what it is to be an American. But her identity as an American is social, tied up with a shared social context in which not everyone agrees with her. And insofar as she sees others as fellow Americans, how they regard the relevance of certain values or traditions bears on her, challenging her conception of what matters. Even with her strong convictions, then, our conservative remains vulnerable to certain others, experiencing herself through them. This leaves those others capable of making a claim on her that goes beyond intellectual exercise. They can challenge her at an existential level. And her strong blaming responses seem to reflect this. (I turn to how my account of the basic concern sheds light on the various reactive attitudes below.)

To briefly summarize, I've suggested that the basic concern has two components. Its object—the first component—is how people regard the relevance of socially salient considerations. The basic concern, that is, is a concern *about* how people regard the relevance of such considerations, which gain their significance against the backdrop of various forms of life that we share with one another. But when we're basically concerned

¹⁷ These theorists, of course, are not the only theorists to subscribe to a view like this. See also Jean Hampton's contribution to Murphy and Hampton 1988; Walker 2006a, 2006b; Hieronymi 2001; and Calhoun 2004, 2016.

about someone's regard for these considerations, their regard also matters to us *in a particular way*. This is the second component of the basic concern. I've suggested that when we're basically concerned about how someone regards the relevance of socially salient considerations, that person's regard matters to us as the regard of a fellow participant in some social context. A fellow participant is someone whose outlook we take to carry a kind of practical authority; their outlook can make a claim on us, it can challenge or affirm our perspective about what should matter in the relevant context. And this is because we take ourselves to share the pertinent social context with them; how we go on is partly a matter of our agreement, our arriving at a shared understanding. Finally, I've suggested that the way in which we take a fellow participant's regard to bear on how we should see things is direct, unmediated. In this respect, the basic concern isn't an intellectual or instrumental concern. The directness with which we take people's regard to bear on our perspective when we have the basic concern is a reflection of the fact that our identities are socially constituted. How we experience ourselves is partly a matter of the social context that we share with the person towards whom we have the basic concern, and so how that person regards the relevance of socially salient considerations can impact how we experience ourselves.

This puts us in a position to distinguish the basic concern from the concerns of the previous section's psychiatrist, parent, and sociologist. All three care about how someone regards the relevance of socially salient considerations, but none seem to take how that person regards the relevance of those considerations to be capable of making a claim on them about how they should regard those considerations. They don't feel the urge to

reckon with the other person's evaluative outlook; they don't take it to bear in this way on how they should regard the relevance of socially salient considerations.

Consider our parent. Perhaps her young child does something mean to a peer at daycare. This might worry the parent, and it might affect how she cares for her child. She might begin taking her child to counseling, for instance, or she might make a point of working on emotion management with him. The parent might also become frustrated with or disappointed in herself; she might become worried about how she is raising her child. None of this, though, is to take her child's regard to challenge her perspective about the relevance of the other child's welfare. The parent doesn't feel the need to reckon with her child's evaluative outlook, at least not in this way. Something similar can be said even if the child acts out against *her*. That is, the child might defy the parent in some way. Perhaps the child keeps bothering her as she is trying to get a form filled out. This might certainly frustrate our parent, and she might even try to bargain with the child. She might, for example, tell the child that he can get ice cream if he sits quietly for the moment. But she doesn't take the boy's insistent nagging, his lack of regard for her interests or time, to make a claim on her, to challenge her perception of the relevance of her interests or time in the context of their relationship. And this, I'd like to suggest, is because the parent doesn't see her child's regard as the regard of a fellow participant. She doesn't take her relationship with him to be founded on a shared understanding of what should matter, how to go on together.

Something similar can be said for the sociologist and psychiatrist. When the sociologist studies some group of people, she may be interested in how they regard

certain socially salient considerations. She might be interested, for example, in how they regard the relevance of the environment or of other people. However, at least when she studies them, she does so from a theoretical remove, trying to scientifically understand why they do what they do. This, though, isn't to engage them as fellow participants. Her stance is analogous to that of a primatologist, for example, studying the interactions of a group of chimpanzees. Likewise, we can imagine a psychiatrist's patient acting out, perhaps screaming mean things about her when they touch upon an uncomfortable topic. And we can imagine the psychiatrist being concerned about this because it bears on understanding her patient, and on diagnosing and treating him. But this isn't yet to see her patient's regard as making a claim on her about how she should regard the relevance of her interests.¹⁸

This helps us get at the distinction between the objective stance and the participant stance. The key difference between these two stances appears to be a matter of different ways in which other people's regard can matter to us. When we're basically concerned about someone's regard we consider them a kind of peer, a fellow participant in some social context, one whose evaluative perspective can make a direct claim on ours. Without this concern, as Stern points out, we lose a key form of mutuality. A sort of one-sidedness thus characterizes the objective stance. Of course, there are various ways in which we can care about others or their regard when we take the objective stance toward them. The sociologist, for example, might welcome or regret what she finds out

¹⁸ The point here isn't that the sociologist and therapist *can't* have the basic concern about their respective subjects. I'm trying to explain, rather, the sense in which they don't have the basic concern. The same applies to the parent.

about how some population regards the relevance of certain considerations; the psychiatrist might care about what her patient's regard presages; and the parent might care about her child's well-being and so might care about the child's regard insofar as it bears on how she should treat or train him.¹⁹ None of this, though, is yet to care about the others' regard as the regard of a fellow participant in some community, relationship, or practice—as someone whose perspective has bearing on what is relevant when one is in these social contexts.

5. Some Virtues of the Account

I believe that the foregoing account of the basic concern has much going for it. In this section, I hope to highlight some of its virtues. I offer three: first, my account of the basic concern fits Strawson's naturalistic picture of moral responsibility; second, my account captures the distinctive sense in which regarding someone as a responsible agent is a form of respect; and, third, my account is a promising way of capturing and unifying the wide class of reactive attitudes. I'll consider each of these virtues in turn.

5.1 Strawson's Naturalistic Picture and the Basic Concern

According to Strawson, the basic concern is supposed to be natural to human beings—"thoroughgoing and deeply rooted," something we are "given with the fact of human society." There are two related dimensions to Strawson's thinking here. First, there is a point about depth: the basic concern is taken to be a deep-seated feature of human psychology or sociality. And second (and perhaps consequently), there is a point

¹⁹ Seth Shabo has helpfully pointed out these various ways of being concerned about other people's behavior or attitudes while taking the objective stance toward that person (Shabo 2012a, 2012b).

about ubiquity: the basic concern is taken to be widespread—perhaps not all human beings have it, but the majority do. These two points are crucial to Strawson’s views about responsibility. They are connected to his doubts about our ability to simply opt out of the participant stance, and they’re also central to his descriptive metaphysical project, which involves laying bare the deep structure of *human* responsibility practices. I believe that the account of the basic concern that I’ve given helps vindicate Strawson’s points about the depth and ubiquity of the basic concern.

Consider, for instance, how we develop into socially and morally competent agents in the first place. From the start, as infants, we depend on other people to develop our sense of how we should regard and interact with the various elements of our environment. This is evident not only in infants’ well-documented imitative drives, but also in such phenomena as social referencing and parental scaffolding.

That infants imitate the behavior of others is well established. Such imitation begins early during the neonatal period with simple facial gestures (Meltzoff and Moore 1977, 1983), and it becomes more complicated (and more intricately interwoven with social development) as infants age. By nine-months infants begin imitating not only simple bodily movements but also how others interact with objects (Meltzoff 1988). And by 18-months infants begin imitating others based on their intentions, not simply what they do (Meltzoff 1995). The emergence of these more complicated forms of imitation dovetails with the development of other important ways that infants use others as sources of information about their environment. By this age, for example, infants begin using the facial and tonal expressions of others to evaluate their situation. In this way, infants use

others' emotional responses as guides for their own behavior—avoiding things that elicit fear, e.g., or investigating things that elicit delighted responses. This social referencing behavior represents a form of “vicarious learning,” through which infants learn how to emotionally respond to various situations (Campos and Sternberg 1981; Repacholi 1998).

But we don't socially develop *merely* by imitating and observing what others do and how they react to situations; other people also actively shape our understanding and social competences from the start. Caretakers, through “scaffolding” behavior, utilize infants' imitative capacities and drives to shape their understanding of various social practices and routines (Bruner 1983). In such scaffolding, a caretaker provides an infant's early attempts at intentional behavior with a certain meaning or structure, allowing the infant to form its behavior into coherent routines that the infant comes to understand and which it can then take up as part of its competence. For instance, an infant might reach for an object it wants, and its caretaker might pick up the object and ask, “Would you like this?” The infant might then respond with some garbled sounds, and the caretaker might say, in response, “Yes? Okay—here you go.” Here, the caretaker is treating the infant as having started and navigated a complex engagement (asking someone for something out of reach), one that the infant doesn't yet have the competence to navigate. But it is through such scaffolding that infants develop the relevant competences (Brazleton and Tronick 1980; Kaye 1982; Trevarthan 1979). As Victoria McGeer puts it, “In the course of normal development, children are thus bootstrapped into regulating their own experiences, feelings, thoughts and actions, not just in concert with others, but in

accordance with the intersubjective norms of a shared psychological practice (McGeer 2007, 153).²⁰

We thus first develop our social and moral competences through interactions with other people. From the beginning, that is, we take our cues from others; this is how we learn what is important in the various social contexts we occupy. It shouldn't be surprising, then, if how other people regard the relevance of socially salient considerations continues to make claims on us, that we continue to seek intersubjective agreement in its own right. This isn't to deny that we eventually develop into competent social and moral agents, capable of independently navigating the various social contexts in which we find ourselves, with our own sense of how to appropriately engage with and regard the things in them. But surely we don't completely slough off the psychological mechanisms at work in our infant minds. It is far more likely that our sociality continues on in a similar way, transformed, perhaps, by our increased awareness of how to go on, but not completely expunged by it (McGeer 2008, 249).

If this is right, then it isn't hard to see why the basic concern, as I understand it, might be deep-seated and ubiquitous. The basic concern appears to be a reflection of human sociality more generally. It seems utterly natural that our social and moral lives would revolve around arriving at shared understandings, given the way we develop into socially and morally competent agents in the first place, by regulating ourselves "in concert with others," as McGeer puts it. Again, this isn't to say that our psychological

²⁰ My understanding in the last two paragraphs owes a great deal to McGeer's artful discussions of the relevant developmental psychological literature (see McGeer 2004, 105-108; 2007, 151-154).

mechanisms don't change as we grow into competent social and moral agents; but it is doubtful that we completely leave behind this aspect of our development.

The depth and ubiquity of the basic concern is further evidenced by simple reflection on adult social and moral life. We care about what other people think, about how they regard the relevance of things that we consider important in the contexts that we share with them. Indeed, it's difficult to imagine our social and moral lives without this. Consider the way we seek affirmation from friends or from the larger community when faced with behavior that contradicts our own sense of how we ought to regard the relevance of something that is socially salient. I might, for instance, ask my roommate whether my resentment of someone is "fair," or, in a happier scenario, I might ask, "Isn't what so-and-so did nice?" These behaviors reflect both the way in which someone can challenge our sense of what's appropriate and the way others can affirm it.

This reliance on others for affirmation is no small matter. Margaret Urban Walker, for instance, has pointed out the important function it can serve in contexts of widespread discrimination: "Most of us," she explains, "regardless of the strength of our convictions, will need some community of moral judgment that provides a reference point and a rallying point for our refusal to accept vicious or demeaning norms and the demeaning treatments they justify. We need this community to keep alive our sense that what is accepted is nonetheless unacceptable" (Walker 2006, 103). Here, Walker is pointing to the difficulty in maintaining a certain normative view in the face of widespread opposition. Maintaining one's normative view in such a situation often

requires a countervailing community of support; it often requires peers, whose evaluative outlooks can serve to affirm one's own outlook.

I believe, then, that the basic concern, at least as I've suggested we should understand it, can plausibly be taken to be a ubiquitous and deep-seated feature of our social and moral lives. It seems to be reflected in how we go on together. Moreover, given the way we develop into socially and morally competent agents in the first place, we have reason to think that the basic concern is tied up with our social psychology more generally.

5.2 The Basic Concern as a Distinctive Form of Respect

A second virtue of my account of the basic concern is that it gives distinctive and intuitive sense to a common suggestion about responsibility. It is common for theorists, particularly of the Strawsonian stripe, to take regarding someone as responsible to be a form of respect. This idea is famously reflected in Strawson's suggestion that our proneness to the reactive attitudes is a way of treating someone as "a member of the moral community" (Strawson 1962, 63; cf. also Stern 1974 and Darwall 2006, 84). But while this Strawsonian suggestion is eminently plausible, it remains unclear precisely in what sense regarding another as responsible is a form of respect, a way of treating someone as a member of the moral community.

The problem is that the notions of "respect" and "membership in the moral community" are ambiguous. When we see a severely mentally ill person's interests as morally mattering in their own right, there seems to be an important sense in which we respect that person as a member of the moral community, even if we don't regard that

person as a responsible agent. We can say something similar about babies and non-human animals, both of whose interests we typically take to have some claim on us, despite our—again—not regarding them as responsible agents. The question, then, is in what sense we treat someone as a member of the moral community, in what sense we respect them, when we're prone to hold that person responsible. The basic concern helps to illuminate this distinctive way of respecting another person.

When we're basically concerned about how someone regards the relevance of socially salient considerations, we take how they see things to matter for how we should see things, for how *we* should regard those considerations. This, I've suggested, is a way of treating that person's regard as authoritative for us, something we must at least reckon with, which can challenge or affirm our evaluative outlooks. The basic concern is thus a way of taking another person's evaluative outlook seriously—it is a way, one can say, of respecting another person as an *active* member of the moral community, or, perhaps more accurately, as a participant in some social context. When we're basically concerned about how someone regards the relevance of socially salient considerations, in other words, we treat that person as someone whose voice should be heard, even if we disagree with what it says. In this way, when we regard someone as responsible—when we're basically concerned about their evaluative outlook—we give that person a say in the way others treat her and in what matters in the context we share with her. This helps to flesh out what Pamela Hieronymi is driving at when she writes, “contrary to the advice we tell schoolchildren, we ought to care what other people think. To not care about what you think is to not care about you. To disregard your evaluation is to disregard you”

(Hieronymi 2001, 549). To disregard your evaluation is to disregard you because it is to exclude you in an important respect from a social context that we *share*.

This highlights an important ethical issue at stake in responsibility. Theorists are often motivated to develop accounts of moral responsibility and morally responsible agency for fear of blaming or punishing people who don't deserve it. But if the foregoing is right, then it makes explicit that we can also harm people in a distinctive way by wrongly *exempting* them from responsibility.²¹

5.3 The Basic Concern and the Wide Class of Reactive Attitudes

The last virtue I wish to highlight concerns the reactive attitudes. The basic concern is conceptually connected to these attitudes. It is, *ex hypothesi*, what leaves us prone to them. Any account of the basic concern, then, should inform those attitudes' nature. Of course, providing a full account of the various reactive attitudes would require far more space than I have here. More modestly, I simply hope to show that my account of the basic concern is a promising way of capturing and unifying this wide class of attitudes. If I'm right about this, then we have reason to favor the account of the basic concern that I've given.

To make my case, I'll focus on a wide range of paradigmatic reactive attitudes: resentment, indignation, guilt, forgiveness, gratitude, hurt feelings, and certain kinds of love. Again, to be clear, I don't intend to give full analyses of these attitudes, nor do I mean to fully address all that has been said about them. I merely mean to suggest that the

²¹ See Morris 1968 for an interesting discussion of this ethical problem in the context of punishment.

basic concern, as I understand it, provides a promising way of understanding the reactive attitudes as a distinctive class of attitudes.

Let's begin with resentment and indignation. Generally, it is recognized that resentment is somehow more "personal" than indignation. Theorists, however, have struggled to unpack this. Most commonly, they distinguish resentment and indignation in terms of their objects. Resentment, it is often said, is a response to someone mistreating *us*, whereas indignation is a response to someone mistreating *another*.²² This, though, doesn't seem adequate. First, it sometimes seems appropriate to feel indignant in response to how someone treats us. I might feel indignant, for example, about a student's lying to me. In such a case, resentment might even seem out of place; it might seem like I'm taking things too personally. Second, it seems like both resentment and indignation are responses to norm violations more generally. For instance, it seems plausible to suggest that one might resent someone for burning a flag, but it doesn't seem any less plausible that one might instead feel indignant about this.

I take the foregoing reflections to suggest that resentment and indignation shouldn't be distinguished in terms of their object. I think both are responses to the same thing: someone's inappropriately regarding the relevance of some socially salient consideration (including, in some cases, us). They aren't distinguished, then, based on

²² This view is often attributed to Strawson; however, his distinction between resentment and indignation seems more complicated. Although he takes indignation to belong to "vicarious" reactive attitudes, he notes that these vicarious reactive attitudes are vicarious in that they are "essentially *capable* of being vicarious" (Strawson 1962, 56). In this respect, Strawson is explicit that "one can feel indignant on one's own account" (56).

what they are a response to. They're distinguished, rather, I'd like to suggest, by the way in which they engage with the significance of someone's inappropriate regard.

Here, I'll follow a number of theorists who have recently characterized blame in terms of "protest." For these theorists, blame protests—challenges, stands up against—an implicit claim in someone's action.²³ This fits nicely with the way I've accounted for the basic concern. When we're basically concerned about how someone regards the relevance of some socially salient consideration, remember, we take how that person regards the relevance of that consideration to be capable of making a claim on us. We thus ascribe a kind of authority to that person's perspective; how they regard things can challenge or affirm how we regard them. In cases of resentment and indignation, we perceive someone as having shown an *inappropriate* level of regard for the relevance of a socially salient consideration. How that person regards things thus defies our expectations and so *challenges* us. Resentment and indignation, we can say, are responses to this challenge. In feeling resentment or indignation, one stands up against and rejects the implicit suggestion that one ought to similarly regard the relevance of the pertinent socially salient considerations. In doing so, one also stands up for the opposite claim, that one ought *not* regard the relevance of those considerations in that way.

How does this help us distinguish resentment and indignation? Resentment, I propose, protests on one's *own* behalf. When we resent someone's inappropriate regard, we stand up at once for ourselves and for our conception of what's appropriate in the relevant context. Indignation, on the other hand, simply protests on behalf of a member

²³ See especially Hieronymi 2001, Smith 2013, and Talbert 2012. For an important forerunner to this view, see Murphy and Hampton 1988.

(or perhaps members) of some relationship, community, or practice. When we feel indignant about someone's inappropriate regard, then, we stand up more generally for what's appropriate in the relevant context, for what people in that social context can expect from others.

Let's return to an example from the previous section, above: the couple who are about to wed. The woman member of this couple, recall, wants to keep her last name, but her fiancé wants her to take his. Let's say that after the woman explains that she feels this name-taking practice is antiquated and sexist, her fiancé suggests that she should still take his last name, to preserve her relationship with his parents. It's easy to imagine the woman resenting this: how dare he suggest that I should sacrifice *my* integrity, *my* identity, *my* name to keep the peace with his unreasonable, behind-the-times parents—I should matter more in the context of our relationship; he should have my back. I take it that this sort of reaction is indicative of resentment. The woman's anger, her resentment, stands up against a suggestion about *her* place in the relationship; she stands up for her sense of self, as it is implicated in this social context. In this respect, the woman takes her fiancé's perceived lack of regard personally; it's a challenge to *her*. Something similar was going on with the conservative American, who we imagined resenting someone for saying "Happy holidays!" This conservative felt personally challenged by such "PC" culture because she felt it challenged her very sense of self. In this respect, her resentment protests on her own behalf; she stands up at once for herself and for her conception of how Americans should regard the relevance of certain traditions or values.

Indignation, on the other hand, has a different emphasis. Our conservative might feel indignant, for example, if she simply takes the “politically correct” person’s regard as an affront to Americans more generally. Similarly, we could imagine the woman from the foregoing couple insisting that her fiancé is not treating her as fiancés should treat their future wives—he should support her, because that is how life-partners should treat each other. What makes such insistence indicative of indignation, rather than resentment, is that even if the woman perceives herself as being inappropriately regarded, she stands up against this inappropriate regard on behalf of herself *as a fiancée*. In this respect, the woman’s response isn’t fundamentally different than how she might feel on behalf of a friend, whose fiancé inappropriately regarded her (the friend). To oversimplify things a bit, if the woman resents her fiancé, she says, in effect, you can’t treat *me* like that, whereas if she feels indignant, she emphasizes that you can’t treat *one* like that.

Both indignation and resentment, then, can be understood as protesting, standing up against, the challenge posed by a fellow participant’s inappropriately regarding the relevance of socially salient considerations. In this respect, both are reflective of the basic concern. The key difference is on whose behalf they stand up against that challenge. Resentment, which is more personal, stands up primarily for oneself, whereas indignation, which is perhaps loftier and less personal, stands up for members of the pertinent relationship, community, or practice more generally. I suspect that this is why indignation is often thought to be more righteous, and sometimes also more self-righteous, than resentment.

It seems, then, that there is a plausible way of conceiving of resentment and indignation that highlights how those attitudes might be undergirded by the basic concern. What about guilt? Guilt, of course, is self-directed. This may seem odd. Do we really have the basic concern toward ourselves? I believe we do. In fact, I suspect that having the basic concern towards others implies having it towards oneself; this is simply part of the basic concern's logic. The basic concern, after all, is a concern about other people's regard understood as the regard of a *fellow* participant. But if someone is our *fellow* participant, our peer, this implies that we understand ourselves as having the same standing, that we are also *their* peers or fellows. This might have been part of what Strawson was getting at when he claimed that the reactive attitudes, including self-directed reactive attitudes, rise and fall together—that “all of these attitudes have common roots in our human nature and our membership of human communities” (Strawson 1962, 58).

I believe, then, that it is crucial that we have the basic concern towards ourselves—that we care about how we regard the relevance of socially salient considerations, and that we care about this out of a recognition that we are fellow participants in a larger social context. How does this help us understand guilt, though? Understanding ourselves as peers in a larger social context, I suggest, changes the significance of what we do and how we regard things. As Hieronymi insightfully notes, when our action makes a “threatening claim”—the kind of claim that calls for protest—that claim “persists in social space” (Hieronymi 2001, 550). When we understand ourselves as existing in a larger social context, as having practical authority with respect

to how others should regard the relevance of socially salient considerations, we understand that we cannot simply expunge what we did or the regard reflected in what we did. Indeed, what we did and how we regarded the relevance of some socially salient consideration doesn't only matter for us; it also matters for our fellow participants, our peers. Guilt, then, drawing on the foregoing protest view of blame, is our way of standing up against our own inappropriate regard. When I'm rude to my friend, for instance, I might feel guilt, and this guilt rejects my own action's suggestion that my friend ought to regard the relevance of her interests in the way that I did in being rude. In this way, I repudiate how I regarded my friend's interests, affirming the importance of those interests in the context of our friendship.

The foregoing also helps to shed light on forgiveness. Forgiveness, of course, is a complicated phenomenon. To simplify things a bit, then, I'll focus here on the paradigmatic case of forgiveness: our forgiving someone for inappropriately regarding *us*. I'll thus forego discussing the interesting possibilities of self-forgiveness and forgiving on behalf of others.

How should we understand forgiveness? Minimally, it involves a change in perspective away from blaming another person. Most theorists understand this change in perspective in terms of a forswearing of hard feelings, particularly of resentment.²⁴ On this approach, when I tell someone that I forgive them, I don't only indicate that there won't be any future recriminations or retribution; I also indicate that I feel differently about them and what they did—I indicate that I've had a "change of heart," to borrow

²⁴ See, among others, Allais 2008, Calhoun 1992, Hieronymi 2001, Martin 2010, and Murphy and Hampton 1988.

Cheshire Calhoun's phrase (1992, 77). In this respect, forgiveness involves either overcoming one's resentment toward another, or else at least coming to change one's view about one's resentment, and so committing to dampening its effects and to trying to overcome it.²⁵

Crucially, however, forgiveness, and particularly its attending change of heart, cannot come about in just any way. Hieronymi, for example, points out that one hasn't genuinely forgiven another if one simply takes a pill that eliminates one's resentment (Hieronymi 2001, 530). Similarly, if the change of heart comes about due to head trauma or amnesia, it doesn't seem as though there is forgiveness. In these scenarios, the change of heart simply happens to the purported forgiver. When we forgive another person, though, it doesn't seem like we can be passive in this way. Rather, forgiveness seems to be a rational activity. As Jeffrie Murphy puts it, forgiveness is "the sort of thing one does for a reason." And this seems to be what differentiates it from merely forgetting, "which may just happen" (Murphy and Hampton, 1988, 15).²⁶

But if forgiveness is the sort of thing one does for a reason, this raises a question: What are the right kinds of reasons to forgive someone? For instance, it doesn't seem I've forgiven someone if I forswear my resentment out of a conviction that they didn't actually perform the relevant action. Similarly, we don't forgive someone if our change

²⁵ One might wonder whether, on my account, forgiveness can involve forswearing one's *indignation*. I suspect that it can. Here, though, I'll focus on resentment.

²⁶ One might worry about the idea that forgiveness is something we *do*. Forgiveness doesn't seem the same as, say, picking up a pencil. Here, I (and I take it Murphy and Hieronymi) have in mind something different from voluntary action. We forgive in the same sense that we paradigmatically form beliefs or emotions. Such doings are a result of our rational activity. They thus seem different from things that simply happen to us.

of heart comes about through the discovery that we were wrong to be offended by what they did in the first place. Accounting for why we forgive people, then, is crucial to giving an account—an “articulate” account—of forgiveness (see Hieronymi 2001). And it is to this end that the basic concern is helpful for understanding the phenomenon.

As a number of theorists have noted, the protest view of blame is particularly well suited for understanding forgiveness.²⁷ Why? To begin, when we forgive someone, our understanding of a number of things has to remain stable. For instance, we must continue to believe that the person we’re forgiving did something wrong, that they were responsible for doing it, and that they did it with the relevant state of mind. What, then, changes when we forgive someone? The protest view of blame helps us answer this question. On the protest view, blame is a way of engaging with the *significance* of what someone did—and such significance is subject to change.

There is thus plausibly an important role for the basic concern to play in accounting for forgiveness. Consider, for instance, Odette, who lies to Akeem. Odette’s lying to Akeem reflects a lack of regard for him. And this might challenge Akeem’s sense of his own importance within the context of their friendship. Akeem thus might resent Odette’s lie and the lack of regard it reveals. After some time, let’s say Odette provides a heartfelt apology—she feels terrible, and she deeply regrets what she did. Why does Odette’s apology give Akeem a reason to forgive her? Odette’s lying suggests that Akeem should regard his interests as less relevant than he does, and Akeem’s protest registers and rejects this suggestion. In feeling guilty and in apologizing, though, Odette,

²⁷ See, e.g., Hieronymi 2001, Martin 2010, and Smith 2013. Here, I draw especially on Hieronymi’s thinking.

whom Akeem still regards as a fellow participant in the relevant context, is repudiating her own action's claim. This diminishes the significance of the suggestion implicit in Odette's initial action, and so it gives Akeem a reason to forgive Odette—he no longer needs to protest what her action suggests about how he should regard the relevance of his own interests; the force of the challenge is diminished.

Of course, more remains to be said; a full account of forgiveness in terms of the basic concern requires further elaboration. Here, however, I only mean to suggest that this points to a promising role that the basic concern plays in our forgiveness. It also helps us to see why apology and guilt might give us such strong reasons for forgiveness.

What, then, about hurt feelings, gratitude, and (some kinds of) love? In what way might the basic concern underlie these reactive attitudes?

Begin with hurt feelings. Not much has been written about this reactive attitude. Nevertheless, I think that we can illuminate hurt feelings in terms of the account of the basic concern that I've provided. First, as Strawson points out, hurt feelings are personal. They are in this way related to resentment: they have to do with how another regards the relevance of certain socially salient considerations, as it implicates *me*. Imagine, for example, that a friend doesn't invite me to a party that she is having. This might plausibly hurt my feelings. It might also provoke resentment. The difference between these responses, I think, has to do with the way I engage with the claim my friend's regard makes on me. As we saw earlier, resentment protests. Resentment is thus a sort of "fight response" (Hieronymi 2001, 545); it involves, as Jean Hampton points out, a sort of defiance (Murphy and Hampton 1988, 58). Hurt feelings, on the other hand, I propose,

are characterized by a kind of acquiescence. My friend's not inviting me to her party still makes a claim on me, challenging how I regard the relevance of me or my interests in the context of our friendship; however, my hurt feelings register and *accede* to this challenge.

Hurt feelings are thus, I think, best understood in terms of the practical authority that we ascribe to people's regard when we have the basic concern toward them. This is the sense in which hurt feelings go beyond mere sadness, frustration, and regret. When my friend doesn't invite me to her party, I might also regret this, feel frustrated by it, or feel sad about it. But if my feelings are hurt, it is because I ascribe a special kind of agency to my friend. How we go on together is partly a matter of our agreement; we in this sense share the relevant social context—she's a fellow participant. In this respect, how she regards what's socially salient can change the actual shape of our relations. My friend's not inviting me to her party hurts, then, not only because it is an unfortunate state of affairs, but also because it is, in an important respect, authorized and legitimated by how she regards me.²⁸

Interestingly, if what I've said about hurt feelings is right, there seems to be a positive analogue to them, what we might call "fellow feelings." Thus, my friend might do something nice for me. Perhaps she checks in with me about how some important

²⁸ One might wonder about the scope of the basic concern, given what I've said about hurt feelings. Don't people sometimes have their feelings hurt by babies or by non-human animals, like cats and dogs? If they do, I suspect that hurt feelings in these cases aren't called for. And this is because babies and non-human animals seem to lack the capacities required to share social contexts, to negotiate the terms of their relationships with other people. I therefore suspect that it is inappropriate to relate to babies and non-human animals through shared understandings of what should matter in the context of our relations. I'll turn to the impropriety of the basic concern in the next chapter. Thanks to Dana Nelkin for pressing me on this point.

event went, or maybe she unexpectedly brings me a sandwich from my favorite deli. Such behavior reflects a level of regard for me within the context of our friendship that might affirm my sense of my own importance in that context, that might in this way affirm my sense of self—we are indeed such good friends! This response, like hurt feelings, accedes to my friend’s regard, except in a positive light.

I bring up fellow feelings in part because I think that they contrast with gratitude. To begin, it is worth distinguishing two senses of gratitude. First, we might be grateful *that* something happened. For instance, I might be grateful that a medical test came back negative, or that the soccer ball went in the net, or even that you fulfilled your promise, perhaps if I doubted that you would. This kind of gratitude, which we can call “propositional gratitude,” can be distinguished from another sense: our gratitude *to* someone *for* doing something. It is this latter sense that will concern me here, for only it seems to be a reactive attitude.²⁹

How, then, does the basic concern underlie this latter sense of gratitude? Such gratitude, like hurt feelings and resentment, is a personal reactive attitude. It responds, roughly, to how someone regards the relevance of our interests in some social context. More specifically, we feel gratitude toward someone for acting in a way that reflects a high level of regard for the relevance of our interests, usually a level of regard that we don’t feel we are entitled to expect from the other person but that we nonetheless appreciate. Given our basic concern, when someone so acts, that person’s action suggests

²⁹ This distinction between two senses of gratitude is common in the philosophical literature. It is originally owed to Walker 1980-81. See also McAleer 2012, from which I take the term “propositional gratitude.”

that we should regard the relevance of our interests similarly. Our gratitude, I propose, serves to positively register this implicit suggestion, but it also serves to deny it. Gratitude, that is, rejects the suggestion that we should similarly regard the relevance of our interests.³⁰

This may seem surprising, but I think there is something intuitive to there being an element of denial in our gratitude. Consider, for instance, how we commonly express gratitude. We often tell people that they “shouldn’t have” done what they did, or that they “didn’t have to,” or that something they did was “too much” or “too nice.” All of these phrases suggest that how someone behaved—how they regarded the relevance of our interests in the pertinent social context—was not called for. That gratitude involves such a rejection or denial also helps to make sense of the significance of ingratitude. We often feel that people *should* feel gratitude—that gratitude is owed. And when people don’t feel gratitude in these circumstances, their lack of gratitude is liable to offend the person who benefited them. Why? Presumably because it suggests that the person thinks that they *should* expect the level of regard the other person showed them. In this respect, ingratitude is offensive because it signals someone’s feeling entitled. Finally, on the flip side of ingratitude, one can show too much gratitude. When one does, this can be offensive insofar as it is distancing. Again, I think the element of denial in gratitude is crucial for understanding this phenomenon. Sometimes gratitude suggests that one doesn’t take oneself to be as close to the person who benefited them as that person might

³⁰ This isn’t to suggest that gratitude has a negative valence, like resentment. The point I mean to make concerns the logic of the emotion, so to speak—the way it engages with the significance of the other person’s regard.

like, and this is because one's gratitude suggests that what that person did wasn't something one should expect from them.

One way of making sense of gratitude, then, in terms of the basic concern, is to note that gratitude serves to deny or reject the claim on us, implicit in someone's behavior, that we should regard the relevance of our interests in a way that is more than we feel we should. Of course, this doesn't finally settle things, but it suggests a plausible way that the basic concern, as I've construed it, might play an important role in gratitude.³¹

This brings us to our final reactive attitude: love. Strawson brings up love when he discusses what the objective stance precludes. He thus takes the participant stance to be necessary for "the sort of love which two adults can be said to feel reciprocally, for each other" (Strawson 1962, 52). I thus hope only to vindicate the thought that the basic concern is a necessary feature of mature, adult loving relationships. And this doesn't seem far-fetched. Such a relationship marks precisely the sort of social context in which the basic concern is operative. It is hard to imagine someone sharing in an adult loving

³¹ One might worry about certain cases in which gratitude seems called for but for something that we *do* expect of the other person. For instance, one might suggest the possibility of someone's (intelligibly) feeling gratitude for a teacher's help or for their parents' efforts in raising them. There are three possibilities I'd like to suggest here. First, it might be that this person is grateful *that* her teacher helped her or *that* her parents showed her so much love. Second, it is possible that, though teachers ought to help students with papers, say, the context in which one exists is one in which most teachers don't do this. In this case, one might feel grateful for a teacher who is helpful, because there is a real sense in which one can't expect it from them. And, finally, I'm tempted to suggest that some cases like the foregoing might be better thought of as cases of what I've called fellow feelings. Indeed, I suspect that some of the puzzles around gratitude might be due to a failure to distinguish what are in fact different positive reactive attitudes. Thanks to Dana Nelkin and Agnieszka Jaworska for suggesting the foregoing cases to me.

relationship without taking the other person's perspective to bear on what is and to what extent things are socially salient in that context.

I take it, then, that the account of the basic concern that I've provided in this chapter can plausibly be taken to underlie the wide class of reactive attitudes. This isn't, of course, to say that the details don't need to be filled in at various points. I merely hope to have shown that the basic concern, as I understand it, provides a promising way forward—it gives us some inkling of how the reactive attitudes are connected as a class. If this is right, then it is a significant mark in favor of the way I've construed the basic concern.

6. Conclusion

In this chapter, I've provided an account of the basic concern. I take this to be an important contribution for our understanding of the moral psychology of responsibility in general. It's also, though, the first step to better understanding morally responsible agency. The basic concern, as we've seen, is constitutive of the stance from which we hold other people responsible. When we treat someone as a responsible agent, then, this presupposes being basically concerned about how they regard the relevance of socially salient considerations. To account for what morally responsible agency requires, then, we must account for what is required for our basic concern to be appropriate. In the next chapter, I discuss what such an account of morally responsible agency might look like, drawing on the account of the basic concern that I've developed here.

Chapter 4

Towards a Concern-Based Account of Morally Responsible Agency

1. Introduction

In my first two chapters, I developed the idea that to understand morally responsible agency we must attend to the nature of the concern—the basic concern—expressed by our practices of holding responsible. This is the concern around which our responsibility practices revolve and which leaves us prone to the various reactive attitudes. In the last chapter, I put forward an account of this concern. In this chapter, then, I'd like to consider the possible implications that my account of the basic concern might have for how we should understand morally responsible agency.

To be clear, my intentions here are exploratory. I hope to sketch the starting points for a theory of morally responsible agency, based on the account of the basic concern that I've provided. I don't, then, take what I say to be the final word on such a theory, nor do I mean to provide a comprehensive account of morally responsible agency. Moreover, I won't argue directly against other, competing accounts. My aim is simply to show the way that the account of the basic concern I've put forward can illuminate how we should understand morally responsible agency. In doing so, I hope to demonstrate the promise of both the concern-based approach and my account of the basic concern.

2. Accounting for Morally Responsible Agency: Two Tasks

We should begin by distinguishing two types of questions that a theory of morally responsible agency might hope to address. The account of the basic concern that I've given has implications for both.

Susan Wolf observes, “we acquire senses of ourselves and (some) others as responsible beings, as beings with a distinctive kind of status in the world, whose relations to each other and to the world at large are different in kind from the relations that nonresponsible beings are capable of forming and having” (Wolf 1990, 5). How should we understand this status, though, of being a responsible agent? As Wolf points out, this is a notoriously difficult question to answer. A morally responsible agent, it is said, is capable of exercising agency with a kind of “depth” or “power” or “significance”; we take such agents “more seriously” in some way, treating them “as persons, and not as objects” (Wolf 1990, 5-6). How should we understand these various descriptions, though, this imagery of depth and significance? Accounting for the distinctive status that we accord morally responsible agents, accounting for the depth or significance that we ascribe to such agents’ actions and attitudes, is the first task that an account of morally responsible agency might undertake. This task is a matter of accounting for what it *means* to be a morally responsible agent.

We can distinguish this first task from a second. Many theorists are motivated to account for morally responsible agency out of concerns about when and whether it is appropriate to regard someone as a morally responsible agent. The second type of question that an account of morally responsible agency might hope to address, then, has to do with what we might call the “ethics” of moral responsibility. Perhaps the most prominent question in this field of inquiry, and the one that I’ll focus on here, concerns

what capacities are required for one to be appropriately regarded as a morally responsible agent.¹

This second field of inquiry, the ethics of responsibility, has perhaps predominated philosophical discussions of morally responsible agency. It is the domain in which debates concerning compatibilism and incompatibilism are had, for example; it is the domain in which the problem of moral luck gets a grip, and in which skeptics raise worries about whether anyone can be appropriately regarded as morally responsible. These discussions, however, it should be emphasized, are dependent on some answer to the first type of question. They depend, in other words, on some conception of what it means to be a morally responsible agent. For example, it is common for theorists to adopt something like Derk Pereboom's "basic desert" sense of responsibility, according to which, roughly, someone is responsible just in case that person would deserve blame for what they did (Pereboom 2014, 2). From this, one can extrapolate a conception of what it means to be a morally responsible agent: a morally responsible agent is the type of agent that can deserve blame. And with this conception, one can intelligibly discuss issues concerning the ethics of responsibility. Under what conditions, one can ask, would it be appropriate to regard someone as the type of agent who can deserve blame?

I doubt that this basic desert sense represents the core notion of what it means to be a morally responsible agent. But I don't mean to challenge Pereboom here, or to argue

¹ Another question that might be raised concerning the ethics of moral responsibility has to do with when one has the standing to hold another person responsible in some way (see, for instance, theorists who have written on the standing to blame: Scanlon 2008, Wallace 2010, Bell 2013, and Todd 2017). We might also wonder about whether there are conditions under which (and, if so, under what conditions) we *ought* to regard someone as a morally responsible agent. Here, I won't address these questions.

against other theorists who adopt similar conceptions of morally responsible agency for the purposes of discussing certain ethical questions.² What is important for my purposes is the distinction that I've drawn between the two types of questions an account of morally responsible agency might hope to answer, along with these questions' logical relation. My strategy in this chapter will be to consider the implications that my account of the basic concern has for how we should answer these questions. I'll start by considering the first question—how we should understand what it means to be a morally responsible agent—before turning to the second, concerning what capacities are required for one to be appropriately regarded as such an agent.

3. What Does It Mean to Be a Morally Responsible Agent?

By understanding morally responsible agency in terms of the basic concern, I'm employing the concern-based Strawsonian strategy. And as I argued in the second chapter, this strategy is particularly well positioned to account for what it means to be a

² As Pamela Hieronymi (Forthcoming) has pointed out, Pereboom's basic desert sense of responsibility is reflective of a broader "merited-consequence" conception of responsibility. Most theorists, Hieronymi notes, adopt some version of a merited-consequence conception, according to which whether someone is responsible for some action or attitude depends on whether that person deserves certain consequences for that action or attitude. Hieronymi specifically cites Pereboom 2014, Rosen 2015, and Zimmerman 2015; but something like the merited-consequence conception is also reflected in many Strawsonian theories of morally responsible agency, particularly those that we canvassed in chapter 2 (e.g., Watson 1987, Wallace 1994, Macnamara 2015b, and Shoemaker 2015). For many Strawsonians, as we saw, a morally responsible agent is the type of agent that can be appropriately responded to with the reactive attitudes. Hieronymi suggests an alternative approach to moral responsibility: rather than focus on merited consequences, focus instead on the way that certain minds can matter to us. In her discussion, though, Hieronymi goes on to adopt a merited-consequence conception, admitting that she "cannot yet articulate the alternative." I take the concern-based strategy to complement Hieronymi's "mattering-minds" proposal, and to provide a further articulation of it.

morally responsible agent. This is because the basic concern constitutes the stance from which we hold people responsible; it represents the concern around which our practices revolve. In this respect, it provides us with a direct way to consider the distinctive kind of agency and the distinctive kind of significance that we ascribe to someone when we hold that person responsible.

As I understand it, the basic concern is a concern about how people regard the relevance of socially salient considerations, considerations that one believes should matter in some social context, given the relationships, communities, or practices that inform that context. When we have the basic concern about how someone regards the relevance of these considerations, we take how that person regards them to be capable of making a claim on us about how we should regard them—that person’s regard can challenge or affirm our own perspective about what should matter to people in the relevant social context. And this is because we care about that person’s regard as the regard of a *fellow participant* in that context. Intersubjective agreement with this person matters in its own right; we see that person as someone whose evaluative outlook carries authority in that how we go on together is a matter of arriving at a shared understanding of the appropriate way to regard the relevance of socially salient considerations.

A crucial aspect of the basic concern, then, is the way it reflects how we share social contexts with (certain) others. In these contexts, what is socially salient is at some level a matter of agreement, of arriving at a shared understanding of how to go on together, and though this doesn’t mean we simply acquiesce to how fellow participants see things, it does mean that we take how they see things seriously.

Given all of this, what does it mean to be a morally responsible agent? A morally responsible agent, on this picture, is a kind of peer, a fellow participant in the foregoing sense. This gives us a prima facie understanding of the kind of status that we accord someone when we regard them as a morally responsible agent, of the kind of significance that we ascribe to their agency. They aren't just obstacles or beneficial features of our environment. Rather, how they exercise their agency matters to us in a particular way—we reckon with it not merely out of instrumental concern, but because we take it to make direct claims on how we should understand what should matter in the context that we share with them. Their regard can thus make a claim on, can challenge or affirm, how we conceptualize the relationships, communities, or practices that inform our shared context.

As Susan Wolf explained, above, morally responsible agency carries with it a distinctive sort of “depth.” The foregoing account of morally responsible agency puts us in a position to unpack this metaphorical language. Consider a point that Herbert Morris (1968) makes about punishment. Morris imagines a system in which we view wrongdoers as pathological, treating them rather than punishing them. In this system, we take the objective stance towards someone when they transgress the norms of the community; we view them as subjects for treatment rather than as morally responsible agents, who can be held responsible for what they do. Morris discusses a number of reasons we might worry about this system of treatment. Of the reasons to worry, though, Morris writes, “perhaps the most frightening of all would be the derogation of all protests to treatment.” He explains:

In a system of punishment a person who has committed a crime may argue that what he did was right. We make him pay the price and we respect his right to

retain the judgment he has made. A conception of pathology precludes this form of respect. (Morris 1968, 487-488)

We can generalize Morris's point. When we view someone as morally responsible, we attribute to that person's regard a kind of depth in the sense that we take them (and thus allow them) to stand behind how they regard things. Taking how that person regards the relevance of socially salient considerations to make a claim on how we should regard those considerations, in other words, means taking that person to be committed to how they regard those considerations; it means allowing that person to stand for how they regard the relevance of certain considerations among fellow participants in the pertinent social context. It thus means treating that person as answerable for their evaluative outlook.

With theorists like T.M. Scanlon and Angela Smith, then, I agree that the core notion of moral responsibility is a kind of answerability.³ I suspect, though, that the sense of answerability at stake in the basic concern is narrower than the sense that Scanlon and Smith might have in mind. As Scanlon originally articulated the idea, answerability is merely rational: one is answerable for some action or attitude insofar as one's action or attitude is a reflection of one's judgments about reasons (Scanlon 1998, 267-272). Angela Smith, who has recently defended this Scanlonian view, has similarly suggested that one is equally answerable for one's mathematical reasoning and for one's moral reasoning. The difference, for Smith, following Scanlon, has to do with the importance of moral violations (see Smith 2015, 116-121). *Pace* Scanlon and Smith, I doubt that the same

³ For the notion of answerability, see Scanlon 1998, ch. 6, and, especially, Angela Smith's work: Smith 2005, 2008, 2012, and 2015. For a sense of how the notion of answerability grows out of Strawson's thinking, see Scanlon 1988, 60-68.

notion of responsibility is operating in, as it were, purely intellectual endeavors as is operating in social and moral life. Although the decisions I make when completing Sudoku are rational—reflections of my judgment—and although one can certainly ask me why I made them, the sense of answerability relevant to social and moral life seems to be a kind of answerability *to* others. So, while there is perhaps a sense in which I’m answerable for decisions that I make as I complete a Sudoku puzzle, I’m not answerable to others for those decisions. A friend might ask me why I made a move, and I might tell her that I made it on a whim, or that it was the only move that could logically be made, or that it is none of her business, something I don’t feel like explaining. While I can answer for my move, then, I’m not answerable to my friend for it.⁴

To better see the difference, we can imagine a case in which completing Sudoku does become something I’m answerable to others about. We can imagine, for example, that my long-time friend and I, when we lived together, used to race to see who could finish Sudoku first. We can further imagine that after many years this friend and I get back together for a sentimental night of Sudoku races. My friend might very well be hurt, and perhaps might come to feel resentment or indignation towards me, if I’m indifferent as we race. In this context, Sudoku racing is conceivably a socially salient consideration. I might feel bad if I’m distracted by work, say—I might sensibly apologize to my friend for my inadequate Sudoku racing.

⁴ Being answerable for something but not to someone for it might be a way to distinguish what theorists have called “attributability” and “accountability.” See Watson 1994, 2011, Fischer and Tognazzini 2011, Shoemaker 2011, 2015, and Smith 2015 for discussions of the different “faces” of responsibility.

I thus agree with Gary Watson (2011), who worries that the foregoing Scanlonian picture of morally responsible agency doesn't conform to the core notion precisely because it doesn't capture the sense in which we're responsible *to* others. The basic concern has a way of spelling out this core notion. As morally responsible agents, we're fellow participants in certain social contexts, and so how we regard the relevance of the socially salient considerations in that context can make a claim on, can challenge or affirm, how other people regard the relevance of those considerations. In this respect, we're answerable to others for how we regard the relevance of those considerations. Insofar as our regard can make a claim on them about how they should regard the relevance of certain socially salient considerations, they can make in turn make a claim on us; they're licensed in taking us to stand for how we regard the relevance of those considerations and, thus, in holding us responsible accordingly. This reciprocity is reflective of the fact that the social context is one premised upon *sharing* an understanding of how to go on together.⁵

Before turning to the upshots that this conception of morally responsible agency has for understanding what capacities such agency requires, it's worth noting one more implication that the basic concern has for how we should think about what it means to be a morally responsible agent. The basic concern is importantly contextual, relative to particular social contexts, informed by particular relationships, communities, and

⁵ I therefore think that theorists who, like Watson, emphasize that blame, or responsibility more generally, has a communicative nature are right. I simply think that this communicative nature is undergirded by something more fundamental, namely, the basic concern. For theorists who subscribe to a communication view of responsibility, see Watson 1987, 2011; Darwall 2006; Shoemaker 2007, 2015; Macnamara 2015a, 2015b; and Fricker 2016.

practices. In this respect, there are different domains of and layers to morally responsible agency. I take this to be a virtue of the account. David Shoemaker has pointed out, for instance, that it seems possible that individuals with mild intellectual disabilities might be responsible only within the context of certain intimate relationships—they might be responsible “to their friends, family, and caregivers” (Shoemaker 2015, 190). This seems eminently plausible. It could be that such agents have the capacity to operate as fellow participants within these intimate social contexts but lack the capacities relevant for being answerable to the larger public. This same thought might also help us better understand some of the complications involved in deciding whether children are morally responsible agents. Within the larger social and political sphere children might be out of their depth; however, perhaps within the context of their relationships with their friends and their family there is more room for their being fellow participants. Indeed, caretakers might even try to cultivate social contexts in which it is easier for children to function as fellow participants, by, for example, making it clear for them what is at stake in the relevant context or by reminding them about the fact that their actions affect other people.⁶

With this conception of morally responsible agency in mind, let’s turn to the second task that an account of such agency might undertake: accounting for the capacities required for one to be appropriately regarded as a morally responsible agent.

⁶ This is reminiscent of a point that Agnieszka Jaworska makes about caring for someone who retains the capacity to value but has impaired means-end reasoning capacities. Caring for such a person might mean trying to help them means-end reason, so that they can best enact their values. See Jaworska 1999, 125-135.

4. When Is It Appropriate to Regard Someone as a Morally Responsible Agent?

When it comes to the ethics of responsibility, perhaps the most pressing question concerns under what conditions it is appropriate to regard someone as a morally responsible agent. On the concern-based strategy, this is a matter of determining under what conditions it is appropriate to have the basic concern about how someone regards the relevance of socially salient considerations. It is, in other words, a matter of determining under what conditions it is appropriate to be invested in someone's regard as the regard of a fellow participant in some social context.

Still, there is some question about how we should understand the notion of "appropriate" here. As Strawson points out, we can deliberately take an objective stance towards others—we can opt out of the basic concern, at least temporarily (Strawson 1962, 52-53). But not all of the reasons we can have for opting out like this seem to bear on whether someone is a responsible agent. For instance, I might have a fraught relationship with my father, who is never shy about expressing his disapproval of me. For the sake of getting through my mother's birthday dinner without incident, then, I might steel myself against him; I might distance myself from him, viewing him with the same eye as I view my tired, 6-month old nephew—a being to be managed, not quarreled with. I might also have reason to take such an objective eye towards my father for the sake of my own mental health, to avoid, in Strawson's words, the strains of involvement. But while expedience and prudence can give me reason, perhaps decisive reason, to abandon the basic concern, while the basic concern might in this sense be inappropriate, these don't seem to be the right kinds of reasons, and this doesn't seem to be the right sense of

propriety, for understanding morally responsible agency and the capacities it requires. My father is a responsible agent. He just happens to be a responsible agent who is difficult to be around.

This is a particular instantiation of a more general problem facing views that, like the concern-based strategy, attempt to understand some evaluative property (e.g., responsibility) in terms of the propriety of certain attitudes (see D'Arms and Jacobson 2000). Indeed, Strawsonian theorists who attempt to understand morally responsible agency in terms of the propriety of the reactive attitudes face the same problem. I might have prudential reason not to respond to my boss's behavior with resentment, and my resentment might be, in this sense, inappropriate; but this says nothing about whether my boss is a morally responsible agent.

How do we resolve this issue? How do we determine the relevant sense of propriety, or determine the kind of reasons that bear on whether someone is a morally responsible agent? The typical approach is to distinguish the right kinds of reasons in terms of the nature of the relevant attitude itself, and I'll follow suit.⁷ Whether it is appropriate to regard someone as a morally responsible agent, then, on the concern-based approach, is a matter of whether the basic concern makes sense on its own terms, given the implicit commitments it has about the nature of its object. We've already seen that when we have the basic concern about how an individual regards the relevance of socially salient considerations, we care about that individual's regard *qua* the regard of a

⁷ Theorists often refer to the relevant kind of normativity here as "fittingness." See D'Arms and Jacobson 2000, 746-747.

fellow participant in some social context. The question, then, is under what conditions it makes sense to see someone as such a fellow participant.

I'd like to propose three conditions that seem relevant. Each has been discussed, to varying degrees, in the literature on morally responsible agency. The three conditions have to do with an agent's evaluative capacities, normative competence, and what has been called "moral sense." I'll consider each in turn.

4.1 Evaluative Capacities

The basic concern is a concern about how someone regards the relevance of certain considerations; it is a concern, in other words, about someone's evaluative outlook, a concern about what they see as worthwhile or important in some social context. When we have the basic concern towards someone, then, we take that person to have certain evaluative capacities. But what sorts of evaluative capacities are these? What sorts of evaluative capacities does the basic concern presume one has?

Agnieszka Jaworska (2016) has helpfully distinguished between two sorts of evaluative capacities.⁸ On the one hand, one might be capable of seeing some consideration as counting in favor of or against certain actions. One demonstrates this capacity when one performs some action because of some consideration, or when some consideration guides one's actions. Having evaluative capacities in this sense isn't very demanding; young children, for instance, seem capable of it (Jaworska 2016, 69-70).

⁸ Jaworska doesn't put her point in terms of evaluative capacities. She distinguishes, rather, between two different senses of someone taking something to be a reason for action, and she saves talk of evaluation for describing the second of these senses. While I'm departing from Jaworska terminologically, I mean my use of "evaluative capacity" to map onto her talk of two different ways one can take something to be a reason.

Still, there is another, more demanding sense in which one can have evaluative capacities. In addition to seeing considerations as counting in favor of or against certain actions, one can see those considerations as good ones to act on. In this respect, one can see the relevant considerations under the guise of the good; one's actions can "express a view of what reasons for acting are *appropriate* in the circumstances at hand" (Jaworska 2016, 70; her emphasis).⁹

Given the nature of the basic concern, given what it means to be a morally responsible agent, it seems apt to regard someone as a morally responsible agent only if that individual has the latter, more demanding evaluative capacity. When we have the basic concern about how someone regards the relevance of socially salient considerations, we take how they regard those considerations to be capable of making a claim on us about how we should regard the relevance of those considerations. And this is because we seek with that person a shared understanding of the appropriate way to regard the relevance of socially salient considerations. It thus seems inapt to have the basic concern toward someone who is incapable of operating under the guise of the good, of committing to the appropriateness of their own outlook. If we take someone's evaluative outlook to be capable of making a claim on us, of challenging or affirming how we perceive what's appropriate, it seems that person should be capable of engaging with the question of the appropriateness of how to regard the relevance of socially salient considerations in the first place.

⁹ The guise of the good serves as a kind of evaluative background. It is not necessarily something that one consciously considers when acting or forming attitudes.

This point is connected to the point I made in the last section, about depth. There, I suggested that the depth of morally responsible agency is tied up with the way in which a morally responsible agent can take a stand about what matters, the way in which a morally responsible agent can commit to their own perspective. Taking a stand like this seems impossible unless one is capable of seeing one's own evaluative outlook as appropriate.

If this is right, if it is only appropriate to regard someone as a morally responsible agent if that individual is capable of operating under the guise of the good, then this gives us reason to think that certain agents aren't morally responsible. Non-human animals, for example, don't seem capable of taking this perspective on their own outlook. Neither, as Jaworska points out, do very young children (Jaworska 2016, 69-70). More interesting still, if morally responsible agency requires these strong evaluative capacities, it seems likely that people with psychopathy aren't appropriately regarded as morally responsible. Drawing on empirical work on the disorder, Gary Watson has recently suggested that people with psychopathy lack a "reflective normative orientation." They appear to lack, in other words, "standards of action, intention, and desire that entail the possibility of conducting oneself badly or incorrectly, and hence serve as the basis for self-criticism and correction" (Watson 2013, 276). If this is right, then it seems to suggest that psychopaths lack the evaluative capacities that being appropriately regarded as a morally responsible agent requires (see Jaworska 2016, 70-74; Nelkin 2016, 58-60).

So, while much of the literature on moral responsibility takes the question of psychopathic agents' responsibility to turn on their moral incapacities (e.g., their lack of

empathy), if our responsibility practices revolve around the basic concern, we have another reason to think that such agents aren't morally responsible.¹⁰ In this way, I join theorists like Antony Duff (1977), Agnieszka Jaworska (2016), and Dana Nelkin (2016), who have all suggested that whether people with psychopathy are morally responsible importantly hangs on whether they have the capacity to value, to operate under the guise of the good.

4.2 Normative Competence

It doesn't seem that having the foregoing evaluative capacities is sufficient for one's being appropriately regarded as a morally responsible agent, however. As children get older, for example, they seem capable of operating under the guise of the good. Nevertheless, we often continue to see them as non-responsible, at least in many contexts. Something similar can be said for people suffering from delusions or for people who are under extremely stressful conditions. Such agents might act under the guise of the good—they might be able to commit to their perspectives as appropriate—but their agency is disrupted at some other level that seems relevant to their responsibility.

A common way that theorists have explained such cases has been by appealing to the notion of normative competence. Having normative competence is, roughly, having competence with the sorts of reasons that should bear on one's decisions or attitudes in

¹⁰ For discussions that focus on the relevance of moral capacities for psychopathic responsibility, see, e.g., Scanlon 1998, 287-290; Fine and Kennett 2004; Levy 2007, 2014; Shoemaker 2007, 2011, 2015; Talbert 2008, 2012; Watson 2011; Nelkin 2015; and Macnamara 2015b.

some particular context.¹¹ In the language of the basic concern, then, normative competence would be a kind of competence with the socially salient considerations relevant to one's decisions or attitudes in some context. Is there reason to think that the basic concern presumes that one possesses competence of this sort? I believe so.

When we have the basic concern towards someone we're invested in that person's regard against the background of social contexts that are premised upon sharing an understanding of how to go on together. This is reflected in the authority that we ascribe to that person's evaluative outlook, the way we take it to be capable of challenging or affirming our own evaluative outlook. Above, I suggested that it doesn't make sense to ascribe such practical authority to someone who isn't capable of evaluating their own perspective as appropriate or inappropriate. Such an agent cannot engage with questions of what's appropriate in some social context, and so it doesn't seem apt to take their outlook to make a claim on us about how *we* should regard the relevance of those considerations. But similarly, I'd like to add, it doesn't make sense to have the basic concern about how someone regards the relevance of socially salient considerations if

¹¹ There are different ways of spelling out the idea of normative competence. Susan Wolf, an early proponent of this notion, puts it in terms of one's "ability cognitively and normatively to recognize and appreciate the world for what it is" (Wolf 1987, 56); Paul Benson talks in terms of one's "ability to criticize courses of action competently by relevant normative standards" (Benson 1987, 469); and R. Jay Wallace writes of "reflective self-control," one's "power to grasp and apply moral reasons, and the power to control or regulate [one's] behavior by the light of such reasons" (Wallace 1994, 157). The general idea that normative competence is relevant to morally responsible agency is widespread. See also, e.g., Fischer and Ravizza 1998; Darwall 2006; and Nelkin 2011. Recently, some theorists have challenged the notion that normative competence is relevant to responsibility. See Scanlon 1998, ch. 6; Smith 2005, 2008; and Talbert 2008, 2012, 2014.

that individual isn't capable of engaging with these questions *competently* in some context.

To be clear, the idea here isn't that morally responsible agents must be successful practical reasoners. There is a sense in which we can say that racists and bigots are incompetently engaging with questions of what's appropriate, but this isn't the sense of competence relevant to determining whether someone is aptly regarded as morally responsible. The relevant competence isn't a matter of whether someone arrives at reasonable answers to normative questions. Rather, the sense of competence that I have in mind involves the more basic capacities required for one to be able to critically engage with questions concerning how one should regard the relevance of socially salient considerations. I believe that there are two levels to such competence.

First, one needs to have a general background understanding of the types of concerns, expectations, vulnerabilities, and dependencies that define the relationships, communities, or practices that are pertinent to the context that one is in. Part of sharing a social context, in other words, is having a general understanding of what sorts of things should matter to people in that context. This is not to say there is no room for disagreement between participants in any given context about what should matter and, especially, the way in which certain considerations should matter. But to intelligibly engage in such contexts, participants need to already share a great deal with other participants; they need to meet some minimal threshold of understanding. Young children, for example, are often out of their element when it comes to much about social and moral life. A child, with her father next to her, might tell a grocery bagger about the

big fight that her parents had the previous night. But this just reflects the fact that she doesn't understand the way considerations like respecting privacy and protecting dignity typically figure into our intimate relationships. Likewise, when a child grabs a stranger's leg, this is reflective of a failure to understand the way bodily integrity and physical space normally factor into social and moral life. These sorts of behaviors are thus reflective of large lacunae in the child's background understanding of what should matter in the social contexts in which she finds herself.

It doesn't seem to make sense to have the basic concern about how someone regards the relevance of socially salient considerations unless that person has a general background understanding of the relationships, communities, or practices that inform the context they are in. Again, the basic concern reflects the way we share in social and moral life, the way we seek shared understandings with others. And if someone lacks a background understanding of the relationships, communities, or practices that inform the context that individual is in, then our interactions with that person cannot yet be premised on arriving at a shared understanding of what's appropriate *in that context*—that person lacks the requisite background. It thus seems inapt to take that person's regard to make a claim on us about how we should perceive what's appropriate in the relevant context.

The first level of normative competence, then, is a matter of whether one has a general grasp of what kinds of things should matter in social and moral life. The second level of competence, on the other hand, has to do with one's ability to use this understanding in the context that one is in. This second level of normative competence is thus a matter of execution; it is a matter of one's ability to use one's general

understanding of the context that one is in to critically engage with how one should regard the relevance of the socially salient considerations in that context. For instance, someone suffering from a delusion might have a general understanding of the sorts of considerations that are relevant to social and moral life. Their delusion, though, might severely interfere with their grasp of reality. It might thus keep them from being able to appreciate what's really at stake in the context that they are in. Likewise, someone under duress, or in a very stressful situation, might have a general understanding of what sorts of things should matter; but, again, their situation might impair their ability to critically engage with the situation they face.¹²

This second level of normative competence, one's executive competence, also seems relevant to whether the basic concern is apt. Even if one has a general understanding of what should matter in some social context, if that person isn't capable of putting that general understanding into action, if how that person regards the relevance of socially salient considerations cannot reflect their general understanding of what's appropriate, then it doesn't seem to make sense to conceive of ourselves as sharing the context with this person in the relevant sense. In other words, because that person is incapable of critically engaging with the situation they face, because they're incapable of bringing their evaluative capacities and background understanding to bear on their practical reasoning, it seems inapt to ascribe the relevant kind of practical authority to their regard, to take how they regard the relevance of socially salient considerations to make a claim on us about how we should regard the relevance of those considerations.

¹² For a helpful discussion of this impairment conception of duress, along with its connection to addiction, see Watson 1999.

And this is because the relevant sort of authority is tied up with our interest in sharing an understanding of how to go on together, and the person doesn't seem to be in a position to engage in this activity with us, at least temporarily.

The foregoing remarks are somewhat tentative; however, it seems to me that normative competence is required for morally responsible agency. The basic concern make sense, I believe, only if the person towards whom we have it has strong evaluative capacities and is in a position to exercise those capacities competently.

4.3 Moral Sense

Finally, I'd like to make some suggestions about one last condition that seems relevant to whether someone is appropriately regarded as a morally responsible agent. Paul Russell (2004) has called this the condition of "moral sense." To be a morally responsible agent, Russell argues, one "must be able to feel and understand moral sentiments or reactive attitudes" (Russell 2004, 293). I believe the concern-based approach to morally responsible agency both helps to illuminate what's involved in moral sense and why it might be required for moral responsibility.

Russell holds that when an agent lacks moral sense "her *sensitivity* to moral considerations is diminished and her motivation to be guided by these considerations is impoverished and limited" (Russell 2004, 294; his emphasis). Such an agent might be able to "recognize and follow moral norms," and she might be able "to anticipate the consequences of any failure to comply with these norms and standards"; however, an agent lacking moral sense is "morally cold," "unaffected by any consideration *of this [moral] kind*." This problem is "deeper," according to Russell, than just one's not

agreeing or accepting “the reactive attitudes that are directed towards [one].” A person without moral sense “cannot even potentially come to share reactive attitudes and feelings because he constitutionally lacks any emotional life of this kind.” There’s thus “no question... of accepting or rejecting the reactive attitudes that others direct at him.” A person without moral sense, Russell insists, “is incapable of any ‘deep assessment’ of himself and others” (Russell 2004, 295).

I think that Russell’s ideas here are complemented by the account of the basic concern that I’ve provided. The basic concern, after all, leaves us prone to the reactive attitudes. And, as I argued at the end of the last chapter, having the basic concern towards someone implies having it towards oneself—to see someone as a *fellow* participant is also to see oneself as a participant. Someone who lacks moral sense, then, I’d like to suggest, building on Russell’s suggestive remarks, lacks the ability to feel and understand the reactive attitudes because she lacks a sense of *herself* as a morally responsible agent, as a fellow participant in social and moral life.

Lacking such a self-conception would have a profound affect on one’s rational capacities. And indeed, part of Russell’s motivation for proffering the condition of moral sense is to help us see what’s missing from views that take morally responsible agency to merely be a matter of rational self-governance or reason-responsiveness.¹³ Someone without moral sense, according to Russell, would conceive of moral (or, we might say, socially salient) considerations merely as external constraints on what he can do. Such an agent would have “an ‘external’ interest in not arousing negative moral sentiments, since

¹³ Particularly, Russell has in mind Wallace 1994 and Fischer and Ravizza 1998.

this will obviously affect the way others treat him,” but he would lack an “internal system.” Consequently, Russell argues, an agent without moral sense “does not really get it.” For an ordinary moral agent, Russell explains, “our capacity to experience and feel moral sentiments, toward ourselves and others, is intimately and inextricably connected with our understanding of the significance of the background moral claims and considerations” (Russell 2004, 296).

Russell doesn't quite spell out the nature of the significance here, the significance that “background moral claims and considerations” come to have for ordinary moral agents, due to their moral sense. If we understand moral sense in terms of our conceiving of ourselves as fellow participants in social and moral life, though, the idea becomes clearer. When we see ourselves as fellow participants, as existing in a larger social space with others, whose regard also bears on us, then how we (and others) regard the relevance of socially salient considerations takes on a new meaning. It becomes public, so to speak; it becomes something for which we're answerable to others. For example, Russell describes the different ways one might respond to being blamed. Someone without moral sense, who operates only with an “external” interest in not arousing negative moral sentiments, would experience negative reactive attitudes as “unpleasant or painful consequences,” something merely to be avoided. Someone with moral sense, on the other hand, can accept being blamed as *appropriate*. This is an interesting phenomenon, because it involves one's endorsing the relevant standards one is being blamed for violating (Russell 2004, 296). Also interesting, though, is the way in which an ordinary moral agent, with moral sense, might respond to being blamed for something that they

don't feel they should be blamed for. One often takes wrongful blame personally, deeply resenting it, or perhaps feeling incredibly hurt by the presumption, depending on the case. This is indicative of feeling challenged in a way that goes far beyond how one would feel if being blamed were of merely instrumental concern.

An agent with moral sense, Russell holds, exercises their capacity for rational self-control differently than an agent without it. The above helps us to see the difference. When we see ourselves as fellow participants, I've suggested, our evaluations take on a new meaning for us. They become public; we become answerable to others for them. How we regard socially salient considerations can thus carry a kind of practical weight, representing a commitment, something we have to be ready to stand up for or to affirm.¹⁴ Thus, as Russell puts it, our "general capacity for rational self-control is... expressed and manifest *through moral sense*" (Russell 2004, 296-297).

I believe, then, that we should understand the notion of moral sense in terms of the basic concern, and particularly the way we conceive of ourselves as fellow participants in social and moral life when we have it. Is Russell right, though, that one is appropriately regarded as a morally responsible agent only if one has moral sense?

¹⁴ Cheshire Calhoun (1995) describes a notion of integrity in this way: acting with integrity means taking one's own judgments about matters seriously, as carrying practical weight for other people. I think this notion of integrity is tied up with one's identity as a fellow participant in social and moral life. Of course, we don't always act with integrity. Still, this doesn't mean one doesn't have moral sense. Indeed, one might feel guilty for one's failings. Interestingly, as an example of someone who acts without integrity in the foregoing sense, Calhoun discusses hypocrites (Calhoun 1995, 258-259). There might thus be a way for the concern-based strategy I'm advocating to speak to the puzzle concerning why hypocrites lose their standing to blame: there is a kind of bad faith in treating others as though they're fellow participants but not in taking oneself to be a participant. (For discussions of hypocrites and the standing to blame, see, e.g., Scanlon 2008 and Wallace 2010.)

I think that he is; the basic concern seems to presume that its object has moral sense. First, as Russell points out, lacking moral sense will surely impair one's development, and so it will likely have an impact on one's normative competence, the range of reasons that one can engage with (Russell 2004; see also McKenna 2012, 82-84). This, though, is an indirect reason for moral sense's relevance for morally responsible agency. More directly, moral sense seems relevant to our ascribing practical authority to someone's evaluative outlook. Earlier, I suggested that ascribing such authority to someone's outlook only makes sense if that person is capable of committing to the appropriateness of their outlook. And one's having moral sense seems tied up with the relevant ability to commit here; it seems tied up with one's evaluative capacities.

More speculatively, we might be able to distinguish a stronger sense of evaluative capacity than Jaworska discusses, and which might be required for the relevant "social" commitment required for morally responsible agency. We could perhaps imagine a rational, asocial agent. Perhaps this agent cares about completing Sudoku. This agent might very well be able to see some consideration as counting in favor of making a particular move, and the agent might even do this under the guise of the good: it is *appropriate* or *correct* to take this consideration as a reason to make this move. But this agent might lack moral sense and the agent might not see herself as answerable to others for anything she does. If this is the case, then perhaps the agent cannot really commit to how she regards anything, at least not in the sense of taking a stand *among others*. And it might be that the basic concern presupposes this latter capacity.

5. Conclusion

In sum, I've suggested that a morally responsible agent is a kind of fellow participant, a peer, someone whose evaluative outlook carries authority for us against the backdrop of social contexts that are driven by shared understandings of how to go on together. Such an agent's regard can challenge or affirm our sense of how we should regard the relevance of socially salient considerations. We thus understand that person's regard as existing in a larger social context, allowing that person to stand for how she sees things. And I've suggested that this might presume at least three kinds of capacities. First, I suggested that the person must be able to operate under the guise of the good; second, I suggested that she must have normative competence within the relevant social context; and, finally, I suggested that she must have moral sense, conceiving of herself as a fellow participant, one whose evaluative outlook carries authority within the social context.

These remarks are meant to be suggestive. There's no doubt, though, that much more work remains to be done. For instance, the requirements I've adumbrated seem intertwined in interesting ways that are worth exploring. Moreover, there are questions about how the theory of morally responsible agency I've sketched connects to other theories, particularly those outside of the Strawsonian line of thought. How does it bear on questions concerning skepticism about moral responsibility? Questions concerning moral luck? Questions concerning compatibilism and incompatibilism? And there are surely more cases to consider. My goal in this chapter, however, hasn't been to settle all of the issues. It's been, rather, to suggest a new way we might engage with them. More

modestly, then, I hope to have at least motivated the potential of the concern-based Strawsonian strategy and the conception of the basic concern that I've put forward.

Conclusion

Endings are sometimes thought of as new beginnings. I'm unsure about whether this is always merited, but it is the sentiment that I'll take in these concluding remarks. I'll begin by briefly summarizing the main ideas that I've developed in this dissertation, before looking ahead, to some questions and projects that remain for another day.

I take myself, in this dissertation, to have put forward a new Strawsonian approach to morally responsible agency. Not only does this "concern-based Strawsonian strategy" better reflect Strawson's thought, in "Freedom and Resentment," but, I've argued, it is also more promising than the reactive attitude-based alternative that has predominated Strawsonian theorizing about morally responsible agency.

In putting forward this concern-based strategy, I've suggested that the basic concern should be at the center of our thinking about responsibility. Rather than focus on the reactive attitudes that characterize our responsibility practices, that is, I've been suggesting that we should focus instead on the way that we're invested in how people exercise their agency when we're prone to hold them responsible for their actions and attitudes; we should focus, in other words, on the concern that leaves us prone to the reactive attitudes in the first place. This is a concern about how people regard the relevance of considerations that we think should matter in the social contexts that we share with them. And when we have it, I've claimed, we take how they regard those considerations to be capable of challenging or affirming our sense of how we should regard those considerations. The basic concern thus ascribes a kind of authority to another person's evaluative outlook, an authority that is reflective of the intrinsic

importance of arriving at a shared understanding of how to go on with that person in the pertinent context. If this is right, then it seems like being a morally responsible agent is being a kind of “fellow participant,” someone whose evaluative outlook carries such authority. And for one to be sensibly regarded as such a fellow participant, I’ve suggested, it seems one must have certain capacities—the capacity to operate under the guise of the good, a normative competence with the relevant social context, and a sense of oneself as a fellow participant in that context.

This, then, is roughly the picture that I’ve painted over the last four chapters. I think that it offers an appealing way to think about the ethics and moral psychology of responsibility. Still, there remains work to be done.

Part of the motivation for this dissertation has been to reorient discussions of responsibility away from the reactive attitudes and toward the basic concern. I believe the tendency to focus on the reactive attitudes has left the basic concern under-theorized. In this respect, I take the account of the basic concern that I’ve proffered to represent a significant contribution to our understanding of the moral psychology of responsibility. Still, I’d like to consider in more detail the relation between this concern and the reactive attitudes to which it, *ex hypothesi*, leaves us prone. I believe that one can spell out in more detail the way the basic concern puts others in a position to emotionally affect us in the way that they do. Moreover, while I gestured at how the basic concern might illuminate the various reactive attitudes, a more detailed account of this should surely be given, not only so that we can better understand the reactive attitudes but also so that we can better understand how the basic concern operates in social and moral life.

I'd also like to consider in more detail the implications that my account of the basic concern has for how we should understand morally responsible agency. In the last chapter, I explored some of the possible implications; but morally responsible agency is a rich subject, one that requires engagement with a variety of cases and complications that I haven't yet been able to work out. To this end, it will also be important to bring my concern-based approach into conversation with non-Strawsonian theorists, along with more traditional worries that motivate discussions of responsibility—worries about determinism, identification, luck, and skepticism. I'm hopeful about this latter project, particularly since the concern-based approach that I've advocated is rooted in "Freedom and Resentment," an essay that was aimed at engaging more traditional debates concerning moral responsibility.

Finally, there are other questions in the ethics of responsibility that loom large and that I haven't really addressed in this dissertation. One particularly important question concerns whether there are conditions under which we *ought* to regard someone as morally responsible. If treating someone as morally responsible is a form of respect, as I've joined others in suggesting, then this gives us reason to think that there are such conditions, that we sometimes go wrong in exempting someone from responsibility. But what are these conditions? And, moreover, how does this imperative to hold people responsible relate to our occasional prudential interest in steeling ourselves against others, adopting the objective stance towards them to avoid the strains of commitment?

Many questions remain, then, surely more than I've mentioned here. But I take these questions to reflect promise, not paucity—a beginning, not an ending.

Bibliography

- Allais, L. (2008.) Wiping the slate clean: The heart of forgiveness. *Philosophy and Public Affairs* 36: 33-68.
- Austin, J.L. (1975.) *How to do things with words*, 2nd edition. J.O. Urmson & M. Sbisà (Eds.). Oxford: Oxford University Press.
- Beglin, D. (2018.) Responsibility, libertarians, and the “facts as we know them”: A concern-based construal of Strawson’s reversal. *Ethics* 128: 612-625.
- Bell, M. (2013.) The standing to blame: A critique. In D.J. Coates & N. Tognazzini (Eds.), *Blame: Its nature and norms* (263-281). Oxford: Oxford University Press.
- Bennett, J. (1980.) Accountability. In Z Van Straaten (Ed.), *Philosophical subjects: Essays presented to P.F. Strawson* (14-47). Oxford: Oxford University Press.
- Benson, P. (1987.) Freedom and value. *Journal of Philosophy* 84: 465-486.
- Brazelton, T.B. and Tronick, E. (1980.) Preverbal communication between mother and infants. In D.R. Olson (Ed.), *The social foundations of language and thought* (299-315). New York: Norton.
- Bruner, J. (1983.) *Child’s talk: Learning to use language*. New York: Norton.
- Butler, J. (2006.) *The works of Bishop Butler*. D.B. White (Ed.). Rochester, NY: University of Rochester Press.
- Calhoun, C. (1992.) Changing one’s heart. *Ethics* 103: 76-96.
- Calhoun, C. (1995.) Standing for something. *The Journal of Philosophy* 92: 235-260.
- Calhoun, C. (2004.) An apology for moral shame. *The Journal of Political Philosophy* 12: 127-146.
- Calhoun, C. (2016.) *Moral aims: Essays on the importance of getting it tight and practicing morality with others*. Oxford: Oxford University Press.
- Campbell, J. (2017.) Strawson’s free will naturalism. *International Journal for the Study of Skepticism* 7: 26-52.

- Campos, J.J. and Sternberg, C.R. (1981.) Perception, appraisal, and emotion: The onset of social referencing. In M.E. Lamb & L.R. Sherrod (Eds.), *Infant social cognition: Empirical and theoretical considerations* (273-309). Hillsdale, NJ: Lawrence Erlbaum Associates.
- Clark, M. (2015.) Nietzsche on free will, causality, and responsibility. In *Nietzsche on ethics and politics* (75-96). Oxford: Oxford University Press.
- Coates, D.J. (2017.) Strawson's modest transcendental argument. *British Journal of the History of Philosophy* 25: 799-822.
- D'Arms, J. and Jacobson, D. (2000.) Sentiment and value. *Ethics* 110: 722-748.
- Darwall, S. (2006.) *The second-person standpoint: Morality, respect, and accountability*. Cambridge, MA: Harvard University Press.
- Dillon, R. (1995.) *Dignity, character, and self-respect*. New York: Routledge.
- Dillon, R. (1997.) Self-respect: Moral, emotional, political. *Ethics* 107: 226-249.
- Duff, A. (1977.) Psychopathy and moral understanding. *American Philosophical Quarterly* 14: 189-200.
- Ekstrom, L. (2000.) *Free will: A philosophical study*. Boulder, CO: Westview Press.
- Fine, C. and Kennett, J. (2004.) Mental impairment, moral understanding, and criminal responsibility: Psychopathy and the purposes of punishment. *International Journal of Law and Psychiatry* 27: 425-443.
- Fischer, J.M. (1994.) *The metaphysics of free will*. Oxford: Blackwell Publishers.
- Fischer, J.M. (2014.) Peter Strawson and the facts of agency. In D. Shoemaker & N. Tognazzini (Eds.), *Oxford studies in agency and responsibility, vol. 2: 'Freedom and Resentment' at 50* (93-118). Oxford: Oxford University Press.
- Fischer, J.M. and Ravizza, M. (1998.) *Responsibility and control: A theory of moral responsibility*. Cambridge, UK: Cambridge University Press.
- Fischer, J.M. and Tognazzini, N. (2011.) The physiognomy of responsibility. *Philosophy and Phenomenological Research* 82: 381-417.
- Frankfurt, H. (1969.) Alternative possibilities and moral responsibility. *Journal of Philosophy* 66: 829-839.

- Frankfurt, H. (1971.) Freedom of the will and the concept of a person. *Journal of Philosophy* 68: 5-20.
- Fricker, M. (2016.) What's the point of blame? A paradigm based explanation. *Nous* 50: 165-183.
- Hampton, J.E. and Murphy, J.G. (1988.) *Forgiveness and mercy*. Cambridge, UK: Cambridge University Press.
- Helm, B. (2001.) *Emotional reason: Deliberation, motivation, and the nature of value*. Cambridge, UK: Cambridge University Press.
- Helm, B. (2010.) *Love, friendship, and the self: Intimacy, identification, and the social nature of persons*. Oxford: Oxford University Press.
- Helm, B. (2014.) Trust as a reactive attitude. In D. Shoemaker & N. Tognazzini (Eds.), *Oxford studies in agency and responsibility, vol. 2: 'Freedom and Resentment' at 50* (187-214). Oxford: Oxford University Press.
- Helm, B. (2017.) *Communities of respect: Grounding responsibility, authority, and dignity*. Oxford: Oxford University Press.
- Hieronymi, P. (2001.) Articulating an uncompromising forgiveness. *Philosophy and Phenomenological Research* 62(3): 529-555.
- Hieronymi, P. (2004.) The force and fairness of blame. *Philosophical Perspectives* 28: 115-148.
- Hieronymi, P. (Forthcoming.) I'll bet you think this blame is about you.
- Hieronymi, P. (Manuscript.) Freedom, resentment, and the metaphysics of morals.
- Jaworska, A. (1999.) Respecting the margins of agency: Alzheimer's patients and the capacity to value. *Philosophy and Public Affairs* 28: 105-138.
- Jaworska, A. (2007.) Caring and internality. *Philosophy and Phenomenological Research* 74(3): 529-568.
- Jaworska, A. (2016.) Holding psychopaths responsible and the guise of the good. In S. M. Liao & C. O'Neil (Eds.), *Current controversies in bioethics* (66-77). New York: Routledge.

- Korsgaard, C. (1996.) Creating the kingdom of ends: Reciprocity and responsibility in personal relations. In *Creating the kingdom of ends*. Cambridge, UK: Cambridge University Press.
- Kane, R. (2016.) Moral responsibility, reactive attitudes, and freedom of will. *Journal of Ethics* 20: 229-246.
- Kaye, K. (1982.) *The mental and social life of babies: How parents create persons*. Chicago: University of Chicago Press.
- Levy, N. (2007.) The responsibility of the psychopath revisited. *Philosophy, Psychiatry, and Psychology* 14: 129-138.
- Levy, N. (2014.) Psychopaths and blame: The argument from content. *Philosophical Psychology* 27: 351-367.
- Macnamara, C. (2015a.) Reactive attitudes as communicative entities. *Philosophy and Phenomenological Research* 90: 546-569.
- Macnamara, C. (2015b.) Blame, communication, and morally responsible agency. In R. Clarke, M. McKenna, & A. Smith (Eds.), *The nature of responsibility: New essays* (211-236). Oxford: Oxford University Press.
- Martin, A. (2010.) Owing up and lowering down: The power of apology. *Journal of Philosophy* 107: 534-553.
- Mason, M. (2003.) Contempt as a moral attitude. *Ethics* 113: 234-272.
- McAleer, S. (2012.) Propositional gratitude. *American Philosophical Quarterly* 49: 55-66.
- McGeer, V. (2004.) The art of good hope. *The Annals of the American Academy* 592: 100-127.
- McGeer, V. (2007.) The regulative dimension of folk psychology. In D.D. Hutto & M. Ratcliffe (Eds.), *Folk psychology re-assessed* (137-156). Dordrecht: Springer: 137-156.
- McGeer, V. (2008.) Trust, hope, and empowerment. *Australasian Journal of Philosophy* 86: 237-254.
- McGeer, V. (2012.) Co-reactive attitudes and the making of moral community. In C. MacKenzie and R. Langdon (eds.), *Emotions, Imagination and Moral Reasoning*. Psychology Press.

- McGeer, V. (2014.) P.F. Strawson's consequentialism. In D. Shoemaker & N. Tognazzini (Eds.), *Oxford studies in agency and responsibility*, vol. 2: 'Freedom and Resentment' at 50 (64-92). Oxford: Oxford University Press.
- McGeer, V. (2015.) Mind-making practices: The social infrastructure of self-knowing agency and responsibility. *Philosophical Explorations* 18: 259-281.
- McKenna, M. (2005.) Where Frankfurt and Strawson meet. *Midwest Studies in Philosophy* 29: 163-180.
- McKenna, M. (2012.) *Conversation and responsibility*. Oxford: Oxford University Press.
- McKenna, M. (2016.) Quality of will, private blame and conversation: Reply to Driver, Shoemaker, and Vargas. *Criminal Law and Philosophy* 10: 243-263.
- Meltzoff, A.N. (1988.) Infant imitation after a 1-week delay: Long term memory for novel acts and multiple stimuli. *Developmental Psychology* 24: 470-476.
- Meltzoff, A.N. (1995.) Understanding the intentions of others: Re-enactment of intended acts by 18-month old children. *Developmental Psychology* 31: 838-850.
- Meltzoff, A.N. and Moore, M.K. (1977.) Imitation of facial and manual gestures by human neonates. *Science* 198: 75-78.
- Meltzoff, A.N. and Moore, M.K. (1983.) Newborn infants imitate adult facial gestures. *Child Development* 54: 702-709.
- Menges, L. (2017.) Grounding responsibility in appropriate blame. *American Philosophical Quarterly* 54: 15-24.
- Morris, H. (1968.) Persons and punishment. *The Monist* 52: 475-501.
- Nelkin, D. (2011.) *Making sense of freedom and responsibility*. Oxford: Oxford University Press.
- Nelkin, D. (2015.) Psychopaths, incorrigible racists, and the faces of responsibility. *Ethics* 125: 357-390.
- Nelkin, D. (2016.) Fine cuts of moral agency: Dissociable deficits in psychopathy and autism. In S. M. Liao & C. O'Neil (Eds.), *Current controversies in bioethics* (47-65). New York: Routledge.

- Nietzsche, F. (1998.) *On the genealogy of morality*. M. Clark and A. J. Swensen (Trans.). Indianapolis, IN: Hackett Publishing Company.
- Pereboom, D. (2014.) *Free will, agency, and meaning in life*. Oxford: Oxford University Press.
- Railton, P. (1984.) Alienation, consequentialism, and the demands of morality. *Philosophy and Public Affairs* 13: 134-171.
- Repacholi, B. (1998.) Infants' use of attentional cues to identify the referent of another person's emotional express. *Developmental Psychology* 34: 1017-1025.
- Rosen, G. (2015.) The alethic conception of moral responsibility. In R. Clarke, M. McKenna, & A. Smith (Eds.), *The nature of responsibility: New essays* (65-88). Oxford: Oxford University Press.
- Russell, P. (1992.) Strawson's way of naturalizing responsibility. *Ethics* 101: 287-302.
- Russell, P. (2004.) Responsibility and the condition of moral sense. *Philosophical Topics* 32: 287-305.
- Scanlon, T.M. (1988.) The significance of choice. In Sterling M. McMurrin (Ed.), *The Tanner lectures on human value*. Salt Lake City: University of Utah Press.
- Scanlon, T.M. (1998.) *What we owe to each other*. Cambridge, MA: Harvard University Press.
- Scanlon, T.M. (2008.) *Moral dimensions: Permissibility, meaning, blame*. Cambridge, MA: Harvard University Press.
- Scanlon, T.M. (2015.) Forms and conditions of responsibility. In R. Clarke, M. McKenna, & A. Smith (Eds.), *The nature of responsibility: New essays* (89-114). Oxford: Oxford University Press: 89-114.
- Shabo, S. (2012a.) Incompatibilism and personal relationships: Another look at Strawson's objective stance. *Australasian Journal of Philosophy* 90: 131-147.
- Shabo, S. (2012b.) Where love and resentment meet: Strawson's interpersonal defense of compatibilism. *Philosophical Review* 121: 95-124.
- Sher, G. (2006.) *In praise of blame*. Oxford: Oxford University Press.
- Shoemaker, D. (2007.) Moral address, moral responsibility, and the boundaries of the moral community. *Ethics* 118(1), 70-108.

- Shoemaker, D. (2011.) Attributability, answerability, and accountability: toward a wider theory of moral responsibility. *Ethics* 121, 602-632.
- Shoemaker, D. (2015). *Responsibility from the margins*. Oxford: Oxford University Press.
- Shoemaker, D. (2017.) Response-dependent responsibility: Or, a funny thing happened on the way to blame. *Philosophical Review* 126: 481-527.
- Smith, A. (2005.) Responsibility for attitudes: Activity and passivity in mental life. *Ethics* 115: 236-271.
- Smith, A. (2007.) On being responsible and holding responsible. *The Journal of Ethics* 11: 465-484.
- Smith, A. (2008.) Control, responsibility, and moral assessment. *Philosophical Studies* 138: 367-392.
- Smith, A. (2012.) Attributability, answerability, and accountability: in defense of a unified account. *Ethics* 122: 575-589.
- Smith, A. (2013.) Moral blame and moral protest. In D.J. Coates & N. Tognazzini (Eds.), *Blame: Its nature and norms* (27-48). Oxford: Oxford University Press.
- Smith, A. (2015.) Responsibility as answerability. *Inquiry* 58: 99-126.
- Stern, L. (1974.) Freedom, blame, and moral community. *Journal of Philosophy* 71: 72-84.
- Strawson, P.F. (1959.) *Individuals: An essay in descriptive metaphysics*. London: Routledge.
- Strawson, P.F. (1961.) Social morality and individual ideal. *Philosophy* 36: 1-17.
- Strawson, P.F. (1962.) Freedom and resentment. *Proceedings of the British Academy* 48: 1-25. Reprinted in J.M. Fischer & M. Ravizza (Eds.), *Perspectives on moral responsibility* (45-66). Ithaca, NY: Cornell University Press.
- Strawson, P.F. (1963.) Carnap's views on constructed systems versus natural languages in analytic philosophy. In P.A. Schlipp (Ed.), *The philosophy of Rudolf Carnap* (503-518). Cambridge: Cambridge University Press.

- Strawson, P.F. (1980.) Replies. In Z. Van Straaten (Ed.), *Philosophical subjects: Essays presented to P.F. Strawson* (260-296). Oxford: Oxford University Press.
- Strawson, P.F. (1985.) *Naturalism and skepticism: Some varieties*. New York, NY: Columbia University Press.
- Strawson, P.F. (1992.) *Analysis and metaphysics: An introduction to philosophy*. Oxford: Oxford University Press.
- Strawson, P.F. (2011a.) Why philosophy? In *Philosophical writings* (222-226). Oxford: Oxford University Press.
- Strawson, P.F. (2011b.) Intellectual autobiography. In *Philosophical writings* (227-247). Oxford: Oxford University Press.
- Talbert, M. (2008.) Blame and responsiveness to moral reasons: Are psychopaths blameworthy? *Pacific Philosophical Quarterly* 89: 516-535.
- Talbert, M. (2012.) Moral competence, moral blame, and protest. *Journal of Ethics* 16, 89-109.
- Talbert, M. (2014.) The significance of psychopathic wrongdoing. In T. Schramme (Ed.), *Being amoral: Psychopathy and moral incapacity* (275-300). Cambridge, MA: MIT Press.
- Todd, P. (2016.) Strawson, moral responsibility, and the “order of explanation”: an intervention. *Ethics* 127: 208-240.
- Todd, P. (2017.) A unified account of the standing to blame. *Nous*: 1-28.
<https://doi.org/10.1111/nous.12215>.
- Trevarthan, C. (1979.) Communication and cooperation in early infancy: a description of primary intersubjectivity. In M. Bullowa (Ed.), *Before speech: The beginning of interpersonal communication* (321-348). Cambridge, UK: Cambridge University Press.
- Walker, A.D.M. (1980-1981.) Gratefulness and gratitude. *Proceedings of the Aristotelian Society* 81: 39-55.
- Walker, M.U. (2006a.) *Moral repair: Reconstructing moral relations after wrongdoing*. Cambridge, UK: Cambridge University Press.
- Walker, M.U. (2006b.) *Moral understandings: A feminist study in ethics, 2nd edition*. Oxford: Oxford University Press.

- Wallace, R. J. (1994.) *Responsibility and the moral sentiments*. Cambridge, MA: Harvard University Press.
- Wallace, R.J. (2010.) Hypocrisy, moral address, and the equal standing of persons. *Philosophy and Public Affairs* 38: 307-341.
- Wallace, R.J. (2011.) Dispassionate opprobrium: On blame and the reactive sentiments. In R.J. Wallace, R. Kumar, & S. Freeman (Eds.), *Reasons and recognition: Essays on the philosophy of T.M. Scanlon* (248-369). Oxford: Oxford University Press.
- Wallace, R.J. (2014.) Emotions and relationships: on a theme from Strawson. In D. Shoemaker & N. Tognazzini (Eds.), *Oxford studies in agency and responsibility, vol. 2: 'Freedom and Resentment' at 50* (119-142). Oxford: Oxford University Press.
- Watson, G. (1975.) Free agency. *Journal of Philosophy* 72: 205-220.
- Watson, G. (1987.) Responsibility and the limits of evil: variations on a Strawsonian theme. In F. Schoeman (Ed.), *Responsibility, character, and the emotions: New essays in moral psychology* (256-286). Cambridge, UK: Cambridge University Press.
- Watson, G. (1996.) Two faces of responsibility. *Philosophical Topics* 24: 227-248.
- Watson, G. (1999.) Excusing addiction. *Law and Philosophy* 18: 589-619.
- Watson, G. (2011.) The trouble with psychopaths. In R.J. Wallace, R. Kumar, & S. Freeman (Eds.), *Reasons and recognition: Essays on the philosophy of T.M. Scanlon* (307-331). Oxford: Oxford University Press.
- Watson, G. (2013.) Psychopathic agency and prudential deficits. *Proceedings of the Aristotelian Society* 113: 269-292.
- Watson, G. (2014.) Peter Strawson on responsibility and sociality. In D. Shoemaker & N. Tognazzini (Eds.), *Oxford studies in agency and responsibility, vol. 2: 'Freedom and Resentment' at 50* (15-32). Oxford: Oxford University Press.
- Watson, G. (2018.) Three faces of responsibility? Comments on responsibility from the margins. *Philosophical Studies* 175: 989-998.
- Williams, B. (1985.) *Ethics and the limits of philosophy*. Cambridge, UK: Cambridge University Press.

- Wittgenstein, L. (1953.) *Philosophical investigations*. New York: The Macmillan Company.
- Wolf, S. (1987.) Sanity and the metaphysics of responsibility. In F. Schoeman (Ed.), *Responsibility, character, and the emotions: New essays in moral psychology* (46-62). Cambridge, UK: Cambridge University Press.
- Wolf, S. (1990.) *Freedom within reason*. Oxford: Oxford University Press.
- Wolf, S. (2015.) Responsibility, moral and otherwise. *Inquiry* 58: 127-142.
- Wonderly, M.L. (2016.) On being attached. *Philosophical Studies* 173: 223-242.
- Zimmerman, M. (2015.) Varieties of moral responsibility. In R. Clarke, M. McKenna, & A. Smith (Eds.), *The nature of responsibility: New essays* (45-64). Oxford: Oxford University Press.