


A Social Practice Account of Responsible Persons

Cheshire Calhoun



OPEN
PRESS
TiU

Commentaries By
Jules Holroyd, Heidi Maibom,
and Gunnar Björnsson

The Descartes Lectures

**A SOCIAL PRACTICE
ACCOUNT OF
RESPONSIBLE PERSONS**

Cheshire Calhoun

Commentaries by Jules Holroyd, Heidi Maibom,
and Gunnar Björnsson

Edited by Miguel Egler and Alfred Archer

© 2024 Miguel Egler and Alfred Archer

Layout design: Proefschrift-aio.nl | Annelies Lips

Cover design: Proefschrift-aio.nl | Guntra Laivacima

ISBN: 9789403771083

DOI: <https://doi.org/10.26116/908v-q016>



The Creative Commons license the author would like to apply for the book is: CC BY-NC-SA 4.0. If this could be placed accompanied by its License icon and the text “This book and all material included in it is licensed under a Creative Commons Attribution-NonCommercial-ShareAlike 4.0 (CC BY-NC-SA 4.0) license. For more information, see <https://www.creativecommons.org/licenses/by-nc-sa/4.0/legalcode>”

The Descartes Lectures

A SOCIAL PRACTICE
ACCOUNT OF
RESPONSIBLE PERSONS

Cheshire Calhoun

Commentaries by Jules Holroyd, Heidi Maibom,
and Gunnar Björnsson

Table of contents

Introduction	6
THE LECTURES	
Chapter 1 Accountability Responsibility	9
1.1 Introduction	9
1.2 Method	13
1.3 Accountability Responsibility	17
1.4 Conclusion	25
Chapter 2 Compliance Responsibility	27
2.1 From Minimal Accountability Responsibility to Robust Compliance Responsibility	27
2.2 Compliance Responsibility	31
2.3 Compliance-Responsibility-Recognizing Attitudes	36
2.4 Compliance-Responsibility-Recognizing Practices	40
2.5 Qualifications	42
2.6 Conclusion	43
Chapter 3 Taking Responsibility	47
3.1 Accommodating Positive Reactive Attitudes	49
3.2 Responsibility Taking and Social Practices	53
3.3 Responsibility Taking	56
3.4 Conclusion	60
Summary Conclusion	63
References	68
COMMENTARIES	
Chapter 4 Non-Idealized Social Practices: Response to Calhoun	75
4.1 Introduction	75
4.2 Rejecting Four Assumptions	77
4.3 Methodology: The Social Practices	80
4.4 Praise and Pressure	90
4.5 Social Practices	99
4.6 References	101

Chapter 5 Being Predictable, Being Trustworthy	105
5.1 What Allows Social Cooperation and Coordination? The Sociopolitical Myth	106
5.2 Social Cooperation and Coordination: The Philosophical Myth	109
5.3 Trustworthiness	118
5.4 Conclusion	121
5.5 References	125
Chapter 6 Responsibility: Expected, Taken, Recognized	129
6.1 Calhoun on Responsible Persons	129
6.2 Do We Assume, by Default, that People Are Responsibility Takers?	130
6.3 Requiring Reasons, Evaluative Autonomy, Balancing Norms	134
6.4 Responsible Behavior, Responsible Persons, and Default Assumptions	139
6.5 What Can We Learn from Positive Reactive Attitudes?	141
6.6 Concluding Remarks	145
6.7 References	147
REPLIES TO COMMENTATORS	
Chapter 7 Björnsson on Taking Responsibility	151
7.1 Preliminaries	151
7.2 Engaging with Björnsson	152
Chapter 8 Holroyd and Non-Idealizing Accounts	159
8.1 Preliminaries	159
8.2 Engaging with Holroyd	161
Chapter 9 Maibom and Social Coordination	167
9.1 Preliminaries	167
9.2 Engaging with Maibom	168
Chapter 10 A Short Note on Gratitude, Praise, and Trust	175
References	177

Introduction

“The Descartes Lectures” is a biennial event at Tilburg University that invites a distinguished philosopher to deliver a series of three lectures, each followed by commentaries from other experts in the field. In 2022, Tilburg University had the honor of hosting Cheshire Calhoun for a series of talks on the important philosophical question of what it means to be a *responsible person*. The commentators for the lectures were Gunnar Björnsson, Jules Holroyd, and Heidi Maibom. This book is a compilation of the material of Calhoun’s lectures, the commentaries by Björnsson, Holroyd, and Maibom, as well as Calhoun’s replies to their critiques.

In her lectures, Calhoun explains that our routine practices of attributing responsibility (to others and to ourselves) challenge the entrenched philosophical view that responsibility can be reduced to accountability to blame. She builds on this discussion to motivate a novel account that aims to do justice to both our normative ideals of what responsible persons *should* be and to our commonsense understanding of what responsible persons *are*. Her main contention is that this approach better captures three key dimensions of being a responsible person: i) the normative requirements of accountability to blame, ii) the expectations we have towards each other, and iii) the disposition of responsible persons to act benevolently, beyond what is expected of them. Calhoun’s lectures thus promise to make headway on thorny theoretical debates about responsibility while urging crucial reflection on the importance of adopting a broader framework to analyze responsible persons that can neatly incorporate both normative and descriptive elements.

The three commentaries on Calhoun’s lectures engage in critical but constructive investigations of the details of Calhoun’s arguments. First, Jules Holroyd endorses Calhoun’s focus on the social practices out of which responsibility practices arise. But she argues that we should take a “non-ideal” approach to these practices, which pays closer attention to the role they play in oppression. Taking this approach, Holroyd argues, results in a more ambivalent, but more realistic, view of the value of responsibility practices.

In the second commentary, Heidi Maibom focuses on Calhoun’s claim that acting in line with our community’s basic social norms is not only something we think people *should* do, but it is also something that we expect that they *will* do. While Calhoun argues for this claim through a detailed look at our social practices, Maibom argues that we can reach a similar conclusion by looking to folk psychology within the philosophy of mind. However, Maibom argues, based on this literature, that we


should place less weight on the role of predictions of other people's behavior and instead focus on the role of effective strategies for dealing with cooperation problems.

The final commentary, by Gunnar Björnsson, examines Calhoun's claim that we have a default assumption that other people will be disposed to take responsibility, meaning that they are disposed to promote good ends in ways that go beyond what is needed for compliance with social norms. Björnsson provides an account of what underlies the normative expectation to promote good ends and argues that this gives us reason to think that we do not generally expect people to be responsibility-takers. Rather, we expect people to comply with the moral norms and one of these norms is that people have a *pro tanto* duty to promote the good. Promoting the good does not go beyond what is normatively expected, then, as it is part of what we expect from others.

Calhoun closes the book by responding to these three commentaries, addressing first Björnsson's critiques, then Holroyd's, and finally Maibom's, noting where she agrees with them and, where she disagrees, before providing additional arguments to support her position.

We are most grateful to Cheshire Calhoun for a rich and insightful set of lectures and to Gunnar Björnsson, Jules Holroyd, and Heidi Maibom for providing such thoughtful and constructive commentaries. The lecture series was funded by the Department of Philosophy at Tilburg University, and we are very grateful for this support. We also wish to thank our fellow organizers of the lecture series, Pilar Lopez-Cantero and Maureen Sie, for their work in making these lectures possible, and Lennart Janssen for providing crucial assistance during the event. Thanks too to all the participants at the Taking Responsibility Workshop, which took place alongside this lecture series, for their contributions to a wonderful three days of philosophical discussion. Finally, we wish to thank Mor Lumbroso, the publishing manager at Open Press TiU, Guus Gijben, our typesetter, and Emma Bolton our copyeditor for their contributions in making this book possible.

Miguel Egler and Alfred Archer



1

ACCOUNTABILITY RESPONSIBILITY

“Morally responsible agency marks a distinctive status; it carries with it a particular sort of social and moral significance. An account of such agency ought to tell us about the nature of this status, this significance.” *David Beglin (2020, 2361)*

1.1 Introduction

Here is the question: “What is a responsible person?” My aims in taking up this question are twofold. First, I aim to loosen the grip that a pervasive view of responsible persons has on philosophers. It’s the view that responsible persons are beings who can be held to account for failing to live up to normative expectations, which is to say, they are those who are liable to blame. It’s not that I think this view is wholly wrong. Accountability, in some sense, is part of being a responsible person, although I’ll be rejecting the identification of accountability with liability to blame and instead suggesting that there’s more to accountability than such liability. More importantly, being accountable is not all there is to being a responsible person. Second, I aim to describe in some detail a more expansive conception of what a responsible person is, of what count as the signature ways of treating persons as responsible (what Strawsonians call “responsibility practices”), and the range of attitudes through which we recognize others as responsible persons. I hope you will find my descriptions utterly familiar from your everyday life with other people. In short, one might say that the goal here is to disrupt entrenched philosophical intuitions about what an account of responsible persons *should* look like in order to capture everyday understandings of what responsible persons *are* like.

The expansive conception of responsible persons that I’ll be developing distinguishes three distinct dimensions of responsible persons. Responsible persons are, first of all,

accountability responsible in the familiar (to philosophers) sense of being capable of living up to normative expectations. Second, they are compliance responsible, which is to say *in fact* disposed to live up to minimal normative expectations, so that many of our normative expectations of responsible persons are also predictive ones. Finally, they are responsibility-takers: they are capable of taking, and are at least sometimes disposed to take, the initiative to do good things that they are not morally required to do. My three lectures will take up each of these dimensions in turn.

Before taking up the first of this trio—accountability responsibility—let me lay more of my cards on the table, starting with my use of the phrase “responsible person.” Philosophers tend to talk either about moral responsibility or about morally responsible agency. I avoid these more familiar terms because of their very strong associations with the project of figuring out whom we can properly make *demands* of to show us a suitable level of regard or respect, and who we can thus *hold accountable* should they fail to do so. I don’t want to bias the investigation into responsible persons from the get-go. I also avoid talk about specifically *moral* expectations, and thus *moral* responsibility for meeting them, in favor of the broader notion of *normative* expectations. Normative expectations cover not only clearly moral ones, but also expectations having to do with etiquette, job responsibilities, the proper ways of doing things such as queuing in line, and so on.

I intend my alternative term, “responsible person,” to draw attention to three points that will be important in what follows: “responsible person” is a *status*; that status is *cross-temporal*; and it is a *default* status in social life.

To have a status is to have a rank in some order of statuses. For example, philosophers are already used to thinking of “moral person” as a status, and take that status to be an especially important and dignified one. For any status, it will be important to ask how we should treat beings with that status, which attitudes toward them are or are not acceptable, and how we can insult them by not treating them or by not having attitudes toward them that befit their status. Kant, for example, thought that contempt was not an attitude we should ever have toward those with the status “moral person.”

The idea that “responsible person” is a status is not entirely unfamiliar. Toddlers and cats, we might say, lack the status “responsible person.” An advantage of focusing on “responsible person” as a status is that it invites us to think about why anyone would *want* this status and would value being recognized and treated as a responsible person and feel aggrieved if they were not. The idea that this status is something we

might prize is not perspicuous when accounts of responsibility emphasize holding other people to account in ways they will no doubt find unpleasant.¹ It's hard to see why anyone would want this status except as a kind of admission price for others' willing interactions.² Moreover, treating our interest in responsibility as primarily an interest in holding others to account shifts attention to the victim's valuable status as a *moral person* who deserves to have their moral personhood recognized in responsible persons' behavior. The agent's own valuable status as a responsible person thus disappears from view.

Nevertheless, the idea that being a responsible person is a valuable status isn't wholly unfamiliar. In Kant's retributivist thinking, to not hold people accountable to the moral and civil law is to fail to recognize and appropriately treat them as having the moral status "person." Thus, the status "responsible person" is valuable because it is a status only moral persons have.

Peter Strawson (2008), by connecting being a responsible person with being viewed from the participant attitude, also suggested that the status "responsible person" is something to be prized. Were we viewed merely from what he calls the "objective attitude," we would be for others merely objects to be managed by pressing the right causal levers. Regarded from the participant attitude, we have for others the distinctive status of being fellow participants in social life and the terminus of interpersonal exchanges. To be regarded and treated as a participant, rather than an object, he thought, just is to be regarded and treated as a responsible person. Although Strawson's emphasis was on our unwillingness—indeed, our likely inability—to give up the idea that others have the status "responsible person," it seems equally true that we ourselves would be unwilling to give up our own status as responsible persons who are fit for interpersonal engagement within social practices.

Assuming that "responsible person" is a valuable status, we can inquire how we should and conventionally do treat persons with that status. Strawson helpfully drew attention to our *responsibility practices*. Although he, and subsequent Strawsonians, took responsibility practices to concern the ways we hold people to account and excuse or temporarily exempt from blame, I will be using "responsibility practices"

¹ Susan Wolf's *Freedom Within Reason* (1993) is an especially notable exception. Throughout, her emphasis is on what she calls the "status of a responsible being" or the "status of responsibility."

² So, Steven Bero (2020), for example, argues that taking responsibility by, for example, expressing contrition and apologizing, is important to us because it assures others that we are eligible for meaningful relationships. And Mark Alfano (2021) argues that we voluntarily accept the burden of potential sanction by taking on responsibilities in order to be seen by others as "worthy partners for future cooperative endeavors" (502).

to refer to the broad spectrum of ways that we treat people as having the status “responsible person,” with particular emphasis on the ways that persons with that status might value being treated and sometimes demand to be treated.

Assuming that “responsible person” is a valuable status, we can also inquire into the attitudes that it is appropriate to have toward responsible persons. Strawson and Strawsonians have focused on what Strawson called “reactive” attitudes, with an overwhelming focus on negative reactive attitudes such as resentment and indignation and an occasional nod to positive attitudes such as gratitude. Because, as I’ll argue, not all of the central attitudes toward responsible persons are *reactive* attitudes that look backward at what has been done, I’ll instead use the descriptor “responsibility-recognizing” attitude. Since the social recognition of persons’ status as responsible persons occurs by way of both the expression of responsibility-recognizing attitudes and responsibility practices of treating them in distinctive ways, failures to extend those attitudes and treatments constitute a distinctive set of status insults.

As a status, being a responsible person is something that one is *cross-temporally*, in the same way one has the statuses “adult,” “middle class,” or “moral person” cross-temporally. Even when we are not responsible for particular actions—when we are excused or temporarily exempted from responsibility—we retain the status “responsible person.” Significantly, the status “responsible person” is a *default* status in social life. Within everyday life, we do not first look for evidence that those we interact with deserve the status “responsible persons.” They are simply presumed to have this status. Of course, we do not presume but instead look for evidence that they are conscientious, dedicated, especially trustworthy people—that is, that they are responsible persons in the sense of having a virtuous character trait. But a status is not a character trait. Strawson thought that it is a central and inextinguishable feature of social life that we adopt the participant attitude toward other people—we see them as the termini of interpersonal interaction, rather than merely as objects to be managed. The participant attitude is a deep and constitutive feature of human social life, one we would be unwilling and likely unable to give up as a general attitude toward others. The participant attitude is thus an attitude of regarding others as having, by default, the status of fellow social participant, which is to say, the status of responsible person. In coming to a conference, in hiring baby-sitters, in sitting peacefully on a train with others, we do not first ask ourselves what evidence there is for regarding all of the people we interact with as social participants with the status “responsible person.” We *assume* they are and are to be treated as such. Where evidence is needed is in supporting our judgments that, in the case of particular individuals, it is a mistake to see and treat them as responsible participants.

So here is a more refined version of my initial question: What does the default, cross-temporal status “responsible person” amount to, and in particular, what competencies ground having that status? How should we treat, and what attitudes should we have toward, individuals who have that default, cross-temporal status? What counts as a status insult to responsible persons?

1.2 Method

In concerning myself with understanding the conception of responsible persons embedded in everyday social life, our responsibility practices, and responsibility-recognizing attitudes, my methodological approach to responsibility falls within the Strawsonian tradition. But it differs in several important ways. I said at the beginning that I intended to avoid expressions like “morally responsible agency” because I do not want to bias the inquiry from the get-go. It makes a methodological and substantive difference whether we begin by focusing on moral responsibility and morally responsible agency, as is typically done, or whether we begin by focusing on the status of responsible person. “Moral responsibility” and “morally responsible agency” naturally invite us to think about what people are responsible, and thus accountable, *for*: What kind of respect or regard do they owe us? What can we demand from them? When are we licensed to react negatively toward—to blame, shun, punish—those who fail to deliver what we normatively expect? The inquiry into moral responsibility thus equally naturally becomes an inquiry into the capacities and features that someone must have if we can properly expect respect or regard from them, can press specific normative demands, and can appropriately hold them responsible and blame them.

If one starts by focusing on responsibility for actual or potential failures to meet normative expectations, it is natural to make the following assumptions about responsible persons, all of which should sound familiar, but all of which I think are mistaken:

1. *The capacities and features of responsible persons are all and only those that license blaming attitudes and holding to account for wrongdoing, absent an acceptable excuse or temporary exemption.*³ If a feature or capacity is not necessary for blaming attitudes to be generally licensed, it is not a feature or

³ Excuses presuppose that a person presently has the requisite capacities but that some factor interferes with their expression in norm-complying action; for example, the person was pushed or in ignorance of relevant facts. Temporary exemption presupposes that the person standardly has the requisite capacities, but they are inoperative due to some factor, for example if the person is suffering a temporary psychotic break.

capacity constitutive of being a responsible being.⁴ The focal contrast, then, is between those who are *liable* (even if sometimes excused or temporarily exempted) and those who are completely *exempted* from liability to blame.

2. *The Strawsonian participant attitude just is the attitude of seeing others as beings of whom we can have normative expectations, and on whom we can make demands*, for a certain kind of regard, goodwill, or respect that recognizes our own status as moral persons.
3. *All responsibility-recognizing attitudes react to blameworthy failures to live up to normative expectations (resentment and indignation) and (in most Strawsonian accounts) to creditworthy exceedings of those expectations (praise, gratitude), and so are properly called “reactive attitudes.”*
4. *Responsibility practices are all and only practices of holding accountable for— or excusing or temporarily exempting from accountability for—actions that fail to meet our normative expectations, or (again, on most accounts) that exceed them.*

So familiar are these four assumptions that you might be mystified as to what else a responsible person could be, or how there could be a *responsibility* practice that is not about holding responsible, or how attitudes fundamentally different from resentment and indignation could be responsibility-recognizing attitudes.

This mystification about what else a responsible person, responsibility practices, and responsibility-recognizing attitudes could be, is not, I think, because there is no other conception of responsible persons, responsibility practices, and responsibility-recognizing attitudes embedded in our everyday interactions with other people within a huge variety of social practices. Indeed, as I suggested at the beginning, one of my goals is to remind you of the much richer, more complexly three-dimensional conception of responsible persons that ordinary people in ordinary social life have. Rather, the sense of mystery, at least for those of us heavily influenced by Strawson, arises from uncritically adopting Strawson’s specific concern about responsibility and his specific methodological approach of looking first to *responsibility* practices rather than to social practices generally.

In his enormously influential essay, “Freedom and Resentment,” Strawson does not take up the general question “What is a responsible person?” Rather, he begins from

⁴ Jules Holroyd calls this the “liability assumption: that to be a responsible agent is to be liable to praise or blameworthiness” (2018, 153).

a very specific concern with responsibility. That concern is with the *propriety* of *holding* people responsible, given uncertainty about what the metaphysical facts are: is determinism true or do people have metaphysically free will? When we hold people responsible, we do unpleasant things to them: blame them, shun them, even jail them. We also demand that they do unpleasant things: feel guilty and remorseful, apologize, make restitution, undertake character reform. We need a justification for doing these things. For the determinist, holding responsible by blaming or punishing must be justified by its utility in altering future behavior. For the libertarian, holding responsible can only be justified if blame or punishment is deserved; and desert depends on the metaphysical freedom of individuals—whether they could have done otherwise. Neither approach seems adequate. The determinist must see individuals as simply objects to be managed by pulling the right causal levers, and thus must exit the participant attitude. Libertarians, while retaining the participant attitude, must rely on an unverifiable, and potentially incoherent, assumption of contra-causal freedom. It is as an intervention into the debate between libertarians and determinists about the propriety of holding responsible that Strawson offers his responsibility-practice account of responsibility. His influential insight was that the accountability of individuals, and practices of holding accountable, do not depend on the truth of any metaphysical view, either determinism or contra-causal freedom. Rather, because practices of holding accountable are essential parts of taking the participant attitude toward others—an attitude that, no matter what the metaphysical facts are, we are unwilling to abandon—we should treat those practices as constitutive of our conception of responsible persons. However, by *starting* from a specific concern with the propriety of holding responsible, we never get a chance to ask, “Is liability to being held accountable all there is to being a responsible person, and are practices connected with holding to account the only responsibility practices?”

Although justifying moral condemnation and punishment is certainly *a* concern we have about responsibility, we need to be open to the thought that individuals’ statuses as responsible persons also matter for reasons *other than* our interest in pressing demands. Mightn’t we also be interested in who can be *predictively* expected to comply with the basic norms that structure social practices? And mightn’t we also be interested in who we can *call on* to volunteer to take on new responsibilities?

In addition to approaching responsibility from the perspective of a specific concern with justifying punitive responses to wrongdoers, Strawson also recommended a specific methodological strategy. We are to begin from the “facts as we know them,” given our experience of social life with others, and then *derive* the conception of responsible persons from those facts. Which facts? If one’s concern is with

the propriety of holding others to account, as Strawson's was and subsequent Strawsonians' has been, the facts are facts about those attitudes and responsibility practices relevant to holding others to account—resentment, indignation, and practices of subjecting to and excusing or exempting from blame. It might seem obvious that these just are the only social facts as we know them that concern responsibility. However, that obviousness is, I suggest, a function not only of the fact that these are responsibility-recognizing attitudes and practices, but also a function of their *salience* to conscious awareness. Resentment and indignation are emotionally felt, and often intensely so. The practices of holding to account—blaming, punishing, demanding apologies, exhorting to better behavior, and so on—are also highly salient to conscious awareness. Such practices involve our *deliberately doing* something, and moreover something that will be unpleasant for the miscreant. Even the acceptance of excuses and extension of temporary exemptions are things typically done after reflection on the evidence.

However, there is no reason to think that, because a set of attitudes and practices are *salient to conscious awareness*, that set is necessarily coextensive with the complete set of responsibility-recognizing attitudes and practices. Suppose that some responsibility-recognizing attitudes and practices are not salient in this way. Perhaps some responsibility-recognizing attitudes are neither felt emotions nor reactions to specific misbehaviors. Perhaps instead they are taken for granted, automatic, and thus unnoticed attitudinal stances. And perhaps some responsibility practices involve *not* doing anything, and in such a way that we don't even notice that there's something we are not doing. This would mean, first, that an account of responsible persons derived only from salient attitudes and practices may be incomplete. Second, and perhaps more worrisomely, the prospects of *deriving* an account of responsible persons from “the facts as we know them” may not work. Instead, we may need to proceed in reverse order, by trying to get a fix on the conception of responsible persons embedded in social life and using that as a guide to identifying responsibility-recognizing attitudes and practices that are not salient to conscious awareness.

Anticipating the second lecture, *the* most pervasive attitude toward responsible persons in everyday life within reasonably well-functioning social practices is basic trust. I don't mean trust in specific individuals—the kind of trust you might decide to invest or find that over the course of repeated interactions you have come to invest. I mean a generalized and default trust that most of the people, largely strangers, that you interact with in assorted everyday social practices—such as sharing trains, shopping at stores, using the library, attending conferences—both know what the

basic normative expectations within those social practices are and will in fact comply with them. This kind of trust, as Annette Baier (1986) observed, is like the air we breathe and is noticed only in its disorienting absence. The responsibility practices that go along with such trust in others' routine compliance with minimal practice norms are exactly what you'd expect—not checking up on people, not taking self-protective measures, not installing surveillance cameras, not insisting on contracts, and so on. We are highly unlikely to notice the things we *don't* do that are nevertheless an important part of treating others as responsible persons.

Given this, I adopt the methodological strategy of starting from social practices generally—not practices of responsibility specifically. We can then ask: what conception of a responsible person is embedded in our social practices? I will argue that it is a complex conception of responsible persons as accountability responsible, compliance responsible, and as responsibility takers.

1.3 Accountability Responsibility

Strawson took accountability responsibility to rest on the capacity to manifest goodwill in one's attitudes and actions. More recently, many have suggested that the basic capacity requisite for accountability responsibility is reasons-responsiveness, which plausibly includes a capacity to understand normative concepts, to detect normatively relevant considerations, and to deliberate on and govern one's actions in light of normatively relevant considerations.⁵ Such a capacity might also rely on emotional capacities like identifying empathy with the effects of one's actions on others.⁶ This, I hope, sounds both familiar and acceptable as a general description of accountability responsible persons' capacities. My aim is not to defend a particular, precisified account, but just to get in view the general conception of an accountability responsible person.

It is, however, important to be clear on what "capacity" means here. Given that the capacities are ones that make one *accountability responsible*, it is very tempting to think that this must be a *realized capacity*. How could anyone be an accountability responsible person who is not *in fact* sensitive to morally relevant considerations, but instead just has a bare, developable but undeveloped capacity? The thought is especially tempting—indeed it seems inevitable, if one accepts the truth of the first assumption I mentioned earlier—that the capacities and features of responsible persons are all and only those that *license* blaming attitudes and holding to account

⁵ See, for example, Wallace 1994; Vargas 2013. The reader should feel free to substitute in their preferred account of the capacities requisite for accountability responsibility.

⁶ See Shoemaker 2020; 2007.

for wrongdoing in the absence of some special excuse. So, let's think about exactly what it might mean to equate being an accountability responsible person with having a realized capacity for reasons-responsiveness. To require that the capacity must be *fully* realized sets the standard for being an accountability responsible person too high. Quite possibly no one meets the idealized standard of being responsive to all normatively relevant considerations. While people might generally be expected to be sensitive to very general and very important considerations (e.g. causing unnecessary pain or humiliation), a large part of our capacity to live up to normative expectations depends on familiarity with specific, local contexts and practices—for example, with the dress norms for different occasions, or the dinner behavior norms in just this family, or standards for ethical medical practice—and this will require quite specific sensitivities.

Avoiding idealizations, one might preserve the connection between being accountability responsible and having realized capacities by adopting Manuel Vargas's (2013) circumstantialist approach. On his view, we should not think of a responsible person—what he calls “morally responsible agency”—as a cross-situational, and by implication cross-temporal, status. If being an accountability responsible person is to license blaming, then the person's capacity for detecting normatively relevant considerations must be *realized*. But since we only have realized capacities with respect to some considerations, in some types of situations, it follows that we sometimes are and sometimes are not accountability responsible persons. We are, that is, not merely excused (which presupposes that we *are* accountability responsible persons⁷), we are totally exempted in some contexts, similar to the way toddlers and cats are exempted in all contexts.

On Vargas's view, as I've said, we must give up the idea that the status “responsible person” is a cross-situational and thus generally *cross-temporal* status. We must also give up the idea that it is a *default* status. After all, whether one is a responsible person or not will vary by context, so we shouldn't presume that individuals are responsible persons across contexts.

Adopting Vargas's circumstantialist view seems exactly right on the condition that one accepts three of the four assumptions I mentioned at the outset of this lecture. Re-stated in simplified form, those three are:

⁷ To be excused is to be in the type of situation to whose normatively relevant features we typically are sensitive—as he puts it, in nearby possible worlds we are reasons-responsive in this type of situation—and our failure to be reasons-responsive on this occasion does not show an absence of good will.

1. The capacities of responsible persons are all and only those that *license* blaming attitudes and holding to account.
2. To regard others as fellow participants is to see them as beings of whom we can make *demands*.
3. Responsibility-recognizing attitudes react to blameworthy failures to live up to normative expectations (resentment and indignation) or to creditworthy exceedings of those expectations (praise, gratitude).

If what having the status “responsible person” gets you is just liability to blame, others’ demands, and resentment, there’s nothing to be lost by sometimes, perhaps often, not having the status “responsible person,” and quite a bit to be gained. Who, after all, wants to have demands pressed upon them and be blamed and resented? More to Vargas’s point, it would be *unfair* to subject people to these things in circumstances where they lacked the realized capacity to understand the relevant norms, to pick up on normatively relevant considerations, and to motivate themselves to comply with those normative expectations they did understand and whose relevance they did pick up on.

The following might seem a perfect example of why being a responsible person should depend on realized capacities, that is, *actual* reasons-responsiveness.

One of our graduate students came from China to live in the U.S. for the first time and study philosophy. By her own account, vast amounts of ordinary everyday normative expectations, including how to properly queue in a grocery store, were entirely unknown to her. Other normative expectations, for example, that one should greet people with a “hello” or display exaggerated (in her view) facial expressions, while known, seemed ridiculous and thus had no motivational grip on her. Often, she could not interpret others’ facial expressions or conversational comments, and thus had no idea whether there was anything of normative significance in them. She couldn’t be constantly asking people “What do you mean by that?” Nor could she rely on having a cultural interpreter with her everywhere she went to act as a norm-explainer and surrogate consideration-detector. And in any case, some normative expectations like the “hello” greeting norm continued to seem ridiculous no matter how their significance was explained (maybe it is thought to be polite, but why have this politeness norm at all?). Lacking many realized capacities for reasons-responsiveness in the U.S., it seems to follow that she frequently does not have the status “responsible person” in the U.S., even though she generally has that status in China.

This was not her view. She took umbrage at philosophers' accounts of responsibility that entailed that, in the U.S. context, she (often) ceased to have the status "responsible person" and thus (often) fell into a category shared by infants, pets, and psychopaths.

What could be her grounds for complaint? It certainly seems inappropriate to hold her to account and blame her for violating normative expectations that were unknown to her or that, because of their alienness, had no motivational grip on her. But if we accept the three assumptions I mentioned and the reasonableness of blame being tied only to realized capacities, what possible grounds for complaint could she have? How could she think she nevertheless had the status responsible person? And what could it mean to treat her as an accountability responsible person if she wasn't eligible for being blamed, having her actions resented, and having demands pressed upon her through all the ways we hold people to account?

Here is how her complaint might be pressed. "Accountability responsible person" is an important and valuable status. True, it sometimes makes us vulnerable to the unpleasant responsibility practices involved in being held to account: being criticized to our faces and behind our backs, ostracized, punished, pressed to apologize, make amends, labor at character reform, and the like. But responsibility practices and responsibility-recognizing attitudes include more than practices of holding to account. Think for example, of workplaces that institute diversity training—what aptly used to be called "sensitivity" training—or that educate hiring committees about how the application and interview process can be structured to reduce the effects of implicit bias. Or think about the enormous amount of work shouldered by members of all sorts of subordinated social groups to educate their wider communities about what the normatively important considerations are, why they are important, and what the new normative expectations should be. These sorts of educative efforts need not, and often don't, presuppose past blameworthiness of their targets, but simply the *developable* capacity to become sensitized to a wider range of normative reasons. In the future, they can do better. In the future, they will be held to account. Thus, one important, and very common, way we treat people as accountability responsible persons is by making the effort to improve their knowledge of norms and sensitivity to normative considerations.

Practices devoted to improving normative competence presume not a realized but a *developable* capacity that makes it worthwhile to engage with others in these ways. To be a responsible person is to be the kind of being who, with help, can *become* sensitive to a wider range of reasons. This is the conception of accountability

responsible persons that Victoria McGeer (2019) develops in presenting her scaffolding-responsiveness conception of responsible persons.⁸ The relevant capacity is, in McGeer's nice phrase, "accordion-like," and expands through exchanges as intelligent people receive and adjust to feedback. She argues that the only capacity for moral reasons-responsiveness required to be an accountability responsible person (to use my term) is "a susceptibility to the scaffolding power of reactive attitudes experienced as a form of moral address" (315). That is, they need only have "whatever it takes to be *sensitizable* to the kind of reasons present at the time of their action, in part by way of the exhortatory effects of (ex-post) reactive scaffolding" (315).

Importantly, I would add, to say that an accountability responsible person is "sensitizable to the kind of reasons present at the time of their action" is not to say that they are sensitizable to what *we* think are the reasons present at the time of action. An accountability responsible person is not necessarily one who *ought* to be receptive to the hortatory and other efforts made to sensitize them to what some "we" take the relevant normative considerations to be. In virtue of having different sensitivities, acquired in very different contexts or through their own reflective efforts, accountability responsible persons sometimes have reason to resist and challenge others' taken-for-granted normative expectations. They can reveal and protest the exclusionary effect of too firmly insisting on norms designed for an insider "we." They can take "us" to task for the uncharitableness of our interpretation of outsiders' lack of conformity to what "we" normatively expect (as, for example, when "we" write them off as backward or misguided). They can press upon us, as we pressed upon them, reasons for acting differently in the future, including reasons for revising particular social norms, such as our "hello" greeting norm.

This last observation points to yet another aspect of accountability responsibility. In narrowly focusing on the practices and attitudes by which we hold others to account, it's easy to overlook that holding accountable is a two-way street. The capacities possessed by others that make holding them to account possible are the very ones that enable them to hold us to account. We treat people as accountability responsible persons in part by being receptive to their own efforts to hold us to account and to scaffold our existing sensitivities. It is an unfortunate feature of hierarchical societies that the moral power to hold to account gets unequally distributed. In Vanessa Carbonell's words, "the marginalized person's moral demand is *ignored, misinterpreted, underestimated, rejected, or silenced*" as a result, for example, of the operation of discrediting stereotypes (2019, 178). Or it may take thousands protesting on behalf of the claims of a single individual or family in order to get responsive uptake (186). "*Claimant*

⁸ She describes this as a skill-based capacity that comes in degrees. See also her 2015.

injustice,” as she says, “occurs when social prejudices or structural inequalities undermine an agent’s ability to engage in felicitous moral address—to make moral claims, to call out wrongdoing, to judge or condemn others for their action, to hold responsible, to seek redress, to blame or punish, or to participate in any of the social practices associated with the participant and vicarious reactive attitudes” (182).

But it is not just social inequality that undermines being treated as a responsible person who can hold others to account. Inability to hold to account is also an unfortunate effect of simply being an outsider to a social practice whose normative expectations are well-entrenched, broadly shared, and have a taken-for-granted legitimacy, so that those who object are written off as simply ignorant. The misfortune in both cases is that responsible persons are treated as less than fully responsible persons on the grounds that they lack what it takes to sensitize others. The misfortune is also that those who could be sensitized by others who are better positioned to detect defects in existing normative expectations fail to profit from those others’ educational efforts.

In sum, the practices of responsibility through which we *treat* individuals as accountability responsible persons are not limited to practices of holding to account and excusing or temporarily exempting. They include trying to improve others’ normative knowledge and sensitivities, giving a hearing to criticisms of and protests against “our” normative expectations, and being open to attempts to press upon us a different set of normative expectations. Thus, we should give up the assumption—number four on my original list—that what count as responsibility practices are all and only practices of holding accountable.

There are also non-blaming responsibility-recognizing attitudes.⁹ These may be expressed with some force. *Disapproval* is the most obvious—for example, disapproval of the normatively ignorant interloper’s misbehavior in a community to which they are not insiders; or disapproval of the normative gaffes occasioned by not being up to date on the nuances of avoiding racist, sexist, or transphobic behavior. There is also the *disappointment* one might feel and express toward those that one hoped might perform better, but who, blamelessly, fail to do so because they’re not “there yet” or because the circumstances are too challenging.¹⁰ There

⁹ Jules Holroyd emphasizes this point in her 2018.

¹⁰ Miranda Fricker contrasts the resentment of disappointment with the resentment of blame. Disappointed resentment falls short of blame, but registers an assessment that the person might have done morally better under difficult circumstances that exempt from blame (2007, 100–105). Adrienne M. Martin similarly discusses both normative hope that individuals will rise above challenging circumstances that would excuse failures to behave well and disappointment when individuals do not rise above those challenges (2014, 129–131).

is also a distinctive *puzzlement* invited by responsible persons' failures to live up to "our" most routine, taken-for-granted normative expectations. One wonders what is going on or what the person means by these failures, and so asks.¹¹ There is also a recognitional attitude appropriately felt toward the *blameworthy* that does not involve blame. It is a kind of hopefulness that might be described as a faith in humanity. By "faith in humanity" I do not have in mind a virtue of optimism, a belief in the inherent goodness and decency of persons, or a willingness to interpret others in the best light.¹² It is rather a stance toward even the most reprobate of allowing that it remains an open possibility, even if remote, that they will improve. They are not to be written off and treated as hopelessly unimprovable.

In sum, we should give up the assumption—number three on my original list—that responsibility-recognizing attitudes always react to blameworthy failures to live up to normative expectations (resentment and indignation) and possibly creditworthy exceedings of those expectations (praise, gratitude).

The Chinese student's complaint is now, I hope, comprehensible: however exempt from blame she might have been, she was not properly exempted from practices of responsibility that treat her as having a developable capacity to expand her sensitivity to normative reasons. Nor was she properly treated as someone unable to critique prevailing normative expectations and hold others to account for defects in their own sensitivities. It is the application of these practices of responsibility that she was still entitled to, and which makes having the status "responsible person" valuable.

The status "accountability responsible person" travels. It travels across contexts where one might be more or less familiar with, or even totally ignorant of, the context-relevant normative considerations. It does so in part because, as McGeer stresses, the status is not grounded (merely) in present realized capacities. It also travels because the requisite capacity is not tied to any particular socio-cultural understanding of the "correct" normative expectations.

¹¹ Think, for example, of how puzzled at what's going on you might have been had you been a subject in the famous Milgram experiment. The sociologist Harold Garfinkel deliberately designed experiments to provoke such puzzlement. In one experiment (1963), his students' task was to behave with the cool politeness of a mere boarder toward family members. While some family members adopted outraged, blaming attitudes (for example, at being called "Mr..." by their own child), others were just baffled and asked if he or she was sick, or had lost their job, or was joking.

¹² For an account of faith in humanity as a virtue, see Preston-Roedder 2013; 2018.

What doesn't travel is liability to blame. I agree with Vargas on this. The status "responsible person" can travel without liability to blame traveling precisely because they are not the same thing. There are two things here: the sometimes merely developable capacities requisite for being an accountability responsible person, and the realized capacities necessary for liability to blame. In short, accountability responsibility is not exclusively a matter of being liable to being held to account, blamed, and resented.¹³

But if to be accountability responsible is not *necessarily* to be liable to—that is, non-exempted from—blame, why call this *accountability* responsibility? Isn't liability to being *held to* account for normative failures, via blame, at the very heart of accountability responsibility? I indicated at the beginning that I aim to work out a conception of responsible persons that is tethered to the everyday conception of responsible persons that shapes our social interactions. In jettisoning the identification of being a responsible person with being liable to blame, haven't I departed from that aim?

I hope what I have said so far goes a considerable distance toward answering that question in the negative. But let me add this: there is a perfectly familiar, everyday notion of *holding* accountable that doesn't involve blame. "Holding accountable" ordinarily means two different things, not a single thing, i.e., preparedness to blame. To illustrate: I hold my students accountable for not plagiarizing. They are familiar with this normative expectation and at least some of the reasons why plagiarizing is bad. Should a student plagiarize, I hold them accountable by automatically failing them. But I also hold my students to performance standards for their papers. They ought, I think, to include thesis sentences, to logically well-order points, to define technical terms, and so on. I down grade them if they do not. But this usually does not involve blame, and the poorer grade is not a punishment. They are just learning what the writing standards are and how to execute them. Both failing students who plagiarize and downgrading students whose writing skills are not up to snuff hold students to normative expectations, but in different ways. As McGeer (2019) emphasizes, one sort of holding looks backward at what could and should have been done differently; the other looks forward to what can and should be done better in the future. The one sense of "holding" presumes a realized capacity for meeting a normative expectation; the other presumes a developable capacity for meeting a normative expectation.

¹³. Jules Holroyd, in her superb essay "Two Ways of Socializing Moral Responsibility" (2018), shows in detail the difficulties of retaining what she calls the liability assumption in an account of responsible agency.


1.4 Conclusion

I've argued that the conception of an accountability responsible person is not reducible to having the capacities that license blame. The capacities that license blame are *realized* capacities. Here I've agreed with Vargas. Sometimes accountability responsible persons have the requisite realized capacities to be liable to blame, sometimes they do not. I've agreed with McGeer that to be an accountability responsible person is to have "whatever it takes" for one's sensitivities to normatively relevant considerations to be *developable*, sometimes through scaffolding interactions with others, sometimes through persons' own critical reflections. And I've observed that realized capacities of accountability responsible persons entail a capacity to challenge *others'* normative understandings and hold *them* to account.

This conception of accountability responsible persons brings into view a wider range of responsibility practices—appropriate ways of treating accountability responsible persons—and a wider range of recognitional attitudes. It does so by not equating being accountability responsible with being liable to blame.

I've argued against any implicit or explicit assumption that accountability responsible persons will be "literate" in local norms. "Illiteracy" may exempt from blame, but not from being treated as having developable capacities or from the standing to challenge local norms. Thus, I've suggested that "accountability responsible person" is not a *social* conception. The status "accountability responsible person" thus travels.

Finally, I hope I've indicated why "accountability responsible person" is a valuable status. While people may not complain about not being resented or blamed, they might well take affront at being written off as uneducable, as enslaved Africans were under American slavery; they might well take affront at being written off as hopelessly irreformable, as criminal offenders may be during and after incarceration; and they might well take affront at being treated as though they lacked what it takes to hold others' accountable or to challenge the legitimacy of prevailing normative expectations.



2

COMPLIANCE RESPONSIBILITY

“A society’s members encounter and know the moral order as perceivedly normal courses of action—familiar scenes of everyday affairs, the world of daily life known in common with others and with others taken for granted.” *Harold Garfinkel (1964, 225)*

2.1 From Minimal Accountability Responsibility to Robust Compliance Responsibility

So far, I’ve been arguing that a familiar view equating accountability responsibility with liability to blame neither captures the full content of accountability responsibility nor the full range of practices and attitudes that recognize individuals as accountability responsible persons. Being accountability responsible, however, is only one of three distinct dimensions of being a responsible person. The second dimension is compliance responsibility.

By way of explaining why the compliance dimension of responsible persons is so important, let me begin by drawing attention both to the *minimalist* nature of conceiving of persons as accountability responsible and to the reasons why that minimalism is often obscured from view. To be an accountability responsible person, I argued, is to have what it takes to either in fact be liable to blame now or to be sensitizable and liable to blame in the future, and, with these, the capacity to blame or sensitize others. In short one has the (at least developable) capacity to live up to others’ normative expectations and to engage others’ developable capacity to live up to one’s own normative expectations.

This is a minimalist conception, first, because having a realized capacity to live up to normative expectations does not mean that one will exercise it—and not just

occasionally not exercise it, but routinely not do so. A person may persistently fail to behave with the most basic common decency or to avoid the most obviously unnecessary harm to others. Perhaps this is due to lazy inattentiveness, or self-indulgent concern with one's own pleasure and interests, or the arrogant thought that norms that apply to others don't apply to oneself, or hostility to some social group. The normatively reprobate do not thereby cease to be accountability responsible persons.

It is not just isolated individuals who may be persistently normatively disappointing. Within some social practices, there may be a widespread disposition to violate practice norms. This is particularly likely within practices where there are significant competitive rewards at stake—rewards of power, wealth, prestige, and the like. Under those conditions, the practice may continue functioning only in virtue of extensive surveillance and penalties that coerce participants' compliance with norms. We also need to keep in mind more serious, pervasive failures to live up to normative expectations that occur under conditions of uncontrollable genocide, terrorist wars, widespread corruption, and the like. Individuals typically remain accountability responsible under these conditions. In short, being an accountability responsible person is compatible with being largely or entirely noncompliant with everyday normative expectations, or compliant only because of effective detection and punitive mechanisms.

Regarding others as accountability responsible persons is, thus, not an *optimistic* stance. A normative expectation of good will is just that—normative. It is not an optimistic expectation that good behavior will normally be forthcoming. But nor is it pessimistic expectation that good behavior will likely not be forthcoming. A *normative* expectation is entirely independent of any beliefs about the likelihood or unlikelihood of others delivering what's normatively expected.

Second, I've been arguing that being an accountability responsible person does not presuppose normative literacy in any particular culture's norms or literacy in a particular practice's norms.¹⁴ Normative illiteracy exempts from liability to blame. It does not exempt from being treated as capable of meeting our normative expectations in the future, or from the accountability responsible person's entitlement to challenge us to change those expectations. Thus, the conception of accountability responsible persons is not a *social* conception in the sense of presupposing prior socialization into "our" norm-governed practices. Accountability responsible persons need neither

¹⁴ My use of the term "normative literacy" is adapted from Barbara Herman's discussion of moral literacy in "Responsibility and Moral Competence" and "Can Virtue be Taught?" in her book *Moral Literacy* (2007).

share normative understandings with us nor care about our norm-governed practices. Consider, by contrast, the following distinctly social conception of accountability responsibility offered by Bennett Helm:

To be accountable is to be answerable to others for upholding the norms of a certain type of community to which they all belong. Such communities ... are communities of respect, in which the members both are jointly committed to some activity or project or way of life (and the norms that define that way of life) and have a kind of standing that each must acknowledge by showing proper respect to others. (2012, 218)

This social conception of accountability responsibility relieves from both accountability, and from being regarded and treated as an accountability responsible person, anyone who is not party to “our” communal, joint commitment to a particular set of norms, either because they simply don’t care about that joint commitment or because, as outsiders, they haven’t been party to it in the first place. I hope the maximalism of this social conception of accountability responsibility strikes you as unacceptable. It exempts from accountability and liability to blame the normatively reprobate. It denies outsiders the status of responsible person, and with that both being treated as having a developable capacity and as having differently sensitized capacities employable in critiquing the normative content of “our” joint commitment. Instead, an appropriate, basic conception of responsible persons as accountability responsible is, and should be, minimalist, relying on a purely normative and non-social construal of “normative expectation,” in order to capture in its net all those who plausibly *are* accountability responsible and should be treated as such.

The minimalism is obscured, I think, by a quite understandable predilection, when we think about accountability responsibility, to follow Strawson in focusing on the “facts as we know them.” What we know are largely facts about what it is like to regard and treat *fellow social participants* in our own, familiar everyday life as accountability responsible. Thus, the envisioned people who serve as paradigms of accountability responsible persons are our friends, family members, neighbors, colleagues, fellow transportation riders, politicians, chat room participants, and so on. They are, as Helm might say, members of our own communities of respect. They have been socialized into roughly the same set of social practices, with their associated norms, as we have. They have also been socialized into many of the same moral norms, including those of common decency, that we have. It’s thus reasonable to assume that those people are already largely sensitized to the same range of considerations that are normatively relevant in particular contexts that we are.

Further, it's reasonable to assume that the considerations they are not yet sensitized to are ones that are relatively easily within their reach. Thus, absent special evidence, it will usually be appropriate to hold them to account through blaming, either because they have likely culpably disappointed normative expectations or because the application of blame effectively speeds up the sensitization process.¹⁵

Because our specific normative demands on those who are fellow social participants are perfectly intelligible to them, expressions of reactive attitudes, such as resentment and indignation, may plausibly be thought of as *communicating* demands that participants are prepared by their social life together to *take up* and appropriately respond to (even if they don't always do so) with self-reactive attitudes such as guilt and remorse, and reparative activities like apologizing and reforming. Among fellow social participants, these reactive attitudes may thus be appropriately thought of as having an essentially call-and-response structure.¹⁶

Finally, and most importantly, it's reasonable to associate normative expectations of particular individuals with predictive expectations about how most will behave. Those who fail to live up to normative expectations fail to do what most others manage to in fact do. Think of familiar examples: the ungrateful gift recipient, the rude cashier, the insulting colleague, the unconscionably racist politician, the deceitful repairman, and so on. In normatively expecting them to behave better than they did, we have in the back of our minds the thought that it's just *normal* to behave better. The resentments of everyday life reflect not only assumptions about how people *ought* to behave, but also assumptions about how people generally do manage to behave. Against this backdrop, talk about "normative expectations" covertly does double duty. It expresses a normative demand and implicitly invokes the normalcy of expecting that people in general *will* live up to that demand, at least with respect to the basics.

There's nothing wrong with this association between normative expectations and general predictive expectations in everyday life. Indeed, as I'll argue in this lecture, I think it's central to our everyday conception of responsible persons that our normative expectations of them are also predictive expectations. But it's important not to build that association into a conception of *accountability* responsible persons. That is a ground-floor conception of responsible persons: if you're going to

^{15.} Victoria McGeer stresses the latter, more instrumental, role of blaming: "On the account I offer, the point and power of blame is to draw wrongdoers into a kind of exchange where they are perforce challenged to exercise their capacities as responsible agents, to reflect on what they have done, whether or not it is legitimate, and if it is not, to take responsibility for what they have done and for what they will do in the future" (2012, 180). Elsewhere she calls this "reactive scaffolding" (2012).

^{16.} See, for example, Darwall 2009; McKenna 2012; Macnamara 2015.

regard and treat others as responsible persons *at all*, you have to think of them as accountability responsible. It's inappropriate at this ground-floor level, however, to bring in more substantive assumptions about them that would make it generally true of responsible persons that they will in fact live up to basic normative expectations, and that those who don't are outliers. And that is because, as I've said, an acceptable conception of accountability responsible persons needs to apply to the normatively reprobate, the stranger to a culture's or practice's norms, the participants in practices that remain viable only through effective detection and penalties, and members of societies that have broken down to the point of genocide or rampant corruption.

If accountability is a minimalist, ground-floor conception of responsible persons, we will need a more robust conception of responsible persons to capture the conception of responsible persons embedded in our everyday social interactions with fellow participants in decently functioning social practices.

2.2 Compliance Responsibility

Daily participation in social practices relies on a more robust conception of responsible persons than mere accountability responsibility. Social life largely proceeds on the assumption that it won't be necessary to press normative demands on others by blaming or by calling upon enforcement mechanisms. Similarly, social life largely proceeds on the assumption that it isn't necessary to explain what the normative expectations are or to make efforts to sensitize others to normatively relevant considerations in the situations in which they find themselves.

Think about hiring a plumber, riding on a bus, going to restaurants, making use of the library, working in your place of employment, receiving postal deliveries, and teaching or taking a class. In all of these, you participate in a variety of established norm-structured practices with which you are deeply familiar and assume others are as well. You are, and you assume the plumber, bus riders, diners, restaurant staff, library patrons, work colleagues, mail person, and class members will, for the most part, behave as they should.

The default presumption in everyday social life is not just that those we interact with *can* learn and *can* comply with normative expectations. The default presumption is that they *have already* learned what's normatively expected and are *disposed to comply* with those expectations—at least the most basic ones. Thus, the *expectations* to which we are prepared to hold people to account are not just *normative* expectations about how people *ought* to behave. They are also *empirical, predictive expectations* about how people largely *will* behave. Put a bit

differently, normatively expected behavior is just *normal* behavior. Thus, the default presumption is that social participants are compliance responsible persons. Any minimally well-functioning social practice operates via this default presumption.

Compliance responsible persons are minimally well-formed social agents, largely fit to be self-directed fellow participants in social life. They do not need to be carefully overseen or coercively managed. This *is* a social conception of responsible persons. The compliance responsibility dimension of responsible persons is acquired through socialization and social experience in which individuals learn at least the most basic norms structuring the practices in which they participate. The normative expectations at issue are the norms structuring established social practices. The norms, whatever the independent legitimacy they might or might not have, are fundamentally *social* norms. We might say that a norm is a social norm when there is a broadly shared understanding of the norm (at least among participants in the practice that is structured by that norm); broadly shared empirical expectations that the norm will generally be followed, and a broadly shared understanding that, regardless of what individuals might personally think about the legitimacy of the norm, it is nevertheless *socially* normative, specifying what ought to be done by social participants and what is socially sanctionable behavior.¹⁷

Since the aim is to capture a *default* presumption of being a responsible person, it's important not to exaggerate the disposition to comply with socially normative expectations. The compliance responsible person is at least a *minimally* well-formed agent, which is to say that they understand and are disposed to comply with *basic* normative expectations operative within the specific social practices (for example, of academia or of bus-riding) in which they participate, as well as basic normative expectations operative within social interaction generally. Basic normative expectations may concern either the *constitutive* norms of practices or non-constitutive norms that are socially understood as matters of *common decency*. Constitutive norms are practice-defining: if one is to count as participating in a particular practice *at all*, one must comply with the practice-defining norms.

^{17.} I have found Cristina Bicchieri's extensive work on social norms very helpful, but have departed from her account of what a social norm is by not making it a necessary condition for a social norm that individuals have a *conditional* preference to follow it. She thinks that what distinguishes personal from social norms is that individuals are willing to follow social norms *only on the conditions* that they expect a sufficient number of others to follow it and believe a sufficient number of others think that it ought to be followed and are willing to apply sanctions. I do not follow her on this point because an aim of childhood and adult socialization, as well as social conversations and negotiations about what the norms should be, is to provide individuals with reasons for thinking the norm is a legitimate one, that is to say, a norm to which adherence should not be conditional in the way Bicchieri describes. See Bicchieri 2017 and 2006.

For example, that customers must pay for items they take out of a store and that cashiers must return change for anything more than the actual price of the item are constitutive norms of shopping. Someone who takes items without paying for them is not shopping, but shoplifting. A clerk who does not return change for payments exceeding the item's cost is not clerking, but expropriating customers' money.

The sociologist Harold Garfinkel (1964) constructed a series of “breaching” experiments both to determine what the constitutive norms of various social practices are, as well as to discover how individuals respond to violations of constitutive rules. In one of his experiments, he directed his students to behave in their family home as a boarder would, using formal modes of address (“Mr...” and “Mrs...”) to parents, avoiding getting personal (for example, not helping themselves to a snack), conducting themselves in a “circumspect and polite fashion,” and speaking only when spoken to. Such behavior violated the constitutive norms for familial interactions. It elicited not only indignation and blaming accusations of being “mean, inconsiderate, selfish, nasty, or impolite,” but also bewilderment about what was going on and efforts to find some explanation for this norm-violating behavior (Was the student sick? Had he lost his job?) (232).

Of course, not all normative expectations concern constitutive norms. There are, for example, many things one ought not to do while grocery shopping, among them: leave one's shopping cart helter-skelter in parking spaces rather than in the provided carrels, replace unwanted items from one's cart on the wrong shelf, comment negatively on what is in other shoppers' carts, allow one's children to run wild in the aisles, talk on the phone while the cashier is trying to check you out. If one does any of these things, one is still shopping, just not in a considerate way. Garfinkel called these non-constitutive practice norms “preferential rules.” Many normative expectations concern preferential rather than constitutive rules. But not all preferential rules have the same status. Some are norms of common decency—the most basic norms that any minimally well-formed social participant should be able to manage to comply with. Which rules have this status may shift over time and there may, obviously, be social disagreement about what is a matter of mere common decency and what a matter of more elevated normative expectation. One possible test of a norm of common decency is to consider how difficult it would be for an average social participant to bring themselves to violate the rule. “Don't negatively comment on other's purchases” might be among those. Other examples of norms of common decency would include not only prohibitions on intentionally and severely harming others, but also the expectation that a “hello” greeting will

be returned rather than ignored and, in the U.S., the expectation that guests will not smoke in one's house unless given permission.¹⁸

The expectation of accountability responsible persons is, as I've said, purely normative and based on the default presumption that individuals have at least a developable, if not realized, capacity to be reasons responsive. Within ongoing social practices, however, the default presumption is that responsible persons have a realized capacity to be reasons-responsive to basic social norms and are disposed to comply with those norms, i.e., that they are compliance responsible. Our normative expectations are thus both normative *and* predictive.

But exactly how should we understand the basis of the disposition to comply with social norms? Since what we are after is a conception of *responsible* persons, the disposition cannot be based solely on desires to avoid social sanction. Effective surveillance to detect norm violation and threat of significantly deterrent penalties may produce widespread, reliable compliance. Such measures, however, respond to the absence of, and substitute for, compliance *responsibility*. But, equally, since what we are after is a *default* conception, the disposition to comply cannot require a voluntary, deliberation-based, reflective decision to accept social norms' authority as both Bennett Helm (in my earlier quote) and Margaret Gilbert suggest. Gilbert, for example, claims that the normativity of a social rule arises from a joint commitment in which the parties to that commitment "*together impose a constraint* on each of the parties with respect to what it is open to him to do, rationally speaking, in the future" (2008, 13). In short, while mere fear of sanction disconnects compliance from the agent's sense of the normativity of social norms, joint commitment idealizes the moral psychology behind the typical social participant's compliance.

The basis of the disposition to comply, I suggest, is instead to be found in what socialization into norm-structured practices provides us. As Cristina Bicchieri (2017) observes, socialization and social experience produce familiarity with possible categorizations of situations—for example, as a family dinner, as highway driving, as a professional exchange—and with scripts for how to proceed, including social norms to be followed, within categorized situations. In learning scripts for how to proceed, one acquires competency within various social practices. Some of that learning is through explicit instruction, as when children are taught to be polite and considerate, or when employees are given job specific codes of conduct or diversity

^{18.} Although not entirely satisfying, I will leave open the question of how, more specifically, to describe the line between constitutive norms and norms of common decency, on the one hand, and more "elevated" norms on the other.

training. Some of what is learned is just absorbed in routine social experiences exposing one to what is normally done or not done by others. Some is acquired in conversations with friends and peers about what is to be done (“Did I tip enough?” “Is one supposed to bring a gift?”) as well as from blaming and praising reactions to one’s normative performance. The result of acquiring social knowledge about how to proceed is practical “know-how.” As James D. Wallace (2008) puts it, “Practices are activities guided by a shared body of practical knowledge—knowledge of how to pursue the activity. Knowledge of how to do something is normative; it is knowledge of how to do it *properly*, knowledge of better rather than worse ways of doing things” (11). Or, to use Barbara Herman’s descriptor, the acquisition of social knowledge about how one ought to proceed constitutes “moral literacy,” or more generally, normative literacy, analogous to becoming fluent in a language.¹⁹

Once a situation is categorized, the norm-encoded script for it is typically automatically activated, and norm compliant behavior flows naturally from one’s grasp of the situation. While norm-conforming behavior can, and sometimes does, follow a deliberative route, much of the time it follows a heuristic one. “According to the heuristic route, norm compliance is an automatic response to situational cues that focus our attention on a particular norm, rather than a conscious decision to give priority to normative considerations” (Bicchieri 2006, 5). Library patrons, for example, need not deliberate about whether to tear useful pages out of reference books, or whether to shout to their friends across the library. Having acquired a script, practical know-how, or normative literacy with respect to library behavior, the minimally well-formed social participant simply does the “to be done” in this situation. Or think back to Garfinkel’s (1964) breaching experiment involving students behaving like boarders in their own family homes. It takes no special motivation—either of self-interest or commitment to a norm—to refer to one’s parents as “mom” or “dad” rather than Mrs. or Mr. Jones. The competent social participant’s “commitment to motivated compliance consists of his grasp of and subscription to the ‘natural facts of life in society’” (236). It was the students’ seeming loss, during the breaching experiment, of their grasp of the natural facts of family life that the parents found so baffling and disturbing. Barbara Herman nicely describes the automatic norm compliance we expect of the normatively literate:

For morality to perform its central function of securing routine action, moral concepts and features of character need to be acquired in the ongoing process of moral education so that a morally literate agent is able to recognize and respond to what is morally salient in

¹⁹ See “Responsibility and Moral Competence” and “Can Virtue be Taught?” in Herman’s 2007.

the routine circumstances she encounters. ... This is, for the most part, nondeliberative. Like the spatially competent agent's ability to move through ordinary doorways without performing any geometric calculations, the morally literate agent moves among persons without the need to think whether she should or could shove them aside, use their body parts for this or that good cause, or tell the truth when asked the time of day. (2013, 31)

To say that socialization equips social participants with what they need in order to routinely comply with social norms is not to say that a “grasp of the ‘natural facts of social life’” is the only thing motivating norm compliance. One consequence of acquiring practice know-how and normative literacy is that, in doing so, one acquires a sensitivity to the fact that violations are likely to be met with such social sanctions as expressed blame, negative gossip, cooled social relations, a spoiled reputation, formal penalties, and the like. Thus, aversion to social penalty has *some* motivational role to play. Similarly, empathy with the effects of norm violation on others (for example, with the bewilderment and hurt the parents of Garfinkel's students felt) has *some* role to play. Similarly, desires to fit in and belong have *some* motivational role to play,²⁰ as do being moved by the thought that others are counting on the agent's compliance (Jones 2017; 1996) and reverence for the community and respect for its norms (Helm 2012). The point here is that the disposition to comply with basic normative expectations is grounded first and foremost in simply grasping how social life works; but that disposition is compatible with a complex of motivations which includes a desire to avoid sanctions, empathetic identification, the need to belong, taking “I am counting on you” as a direct reason for action, reverence for the community and its norms, and even reflective acceptance of the legitimacy of a normative expectation. Any of these motivations, alone or together, may be more or less centrally motivating on a particular occasion.

2.3 Compliance-Responsibility-Recognizing Attitudes

So far, I have focused on the normative-cum-predictive expectations associated with regarding others as compliance responsible persons, as well as on the nature of being compliance responsible. But, just as there are distinctive responsibility-recognizing attitudes connected with accountability responsibility, so there is a distinctive responsibility-recognizing attitude associated with compliance responsibility. That attitude is not *reactive* to what has been done, and thus the Strawsonian terminology

²⁰ Victor Fernandez Castro and Elisabeth Pacherie (2021) argue that the need to belong explains why we care about others' opinions of us and are later motivated by social expectations, reputational considerations, and such social emotions as guilt, embarrassment, and pride.

of “reactive attitudes” is not appropriate. The central compliance-responsibility-recognizing attitude is prospective, basic (or default) trust.

Philosophers have for the most part focused their attention on the non-default form of trust, where trust is a matter of conscious or non-conscious *accepted* vulnerability to harm from particular others. Such trust is *invested*, wisely or unwisely, in particular people and is typically the *outcome* of social experience with those persons.²¹ Non-default, selective trust emerges under conditions where norm compliance can’t be presumed and thus where there is some question about whether others will or will not live up to normative expectations. This may be because the norm in question is not a basic norm, or because a practice is not well-functioning (instead depending on coercive management of participants), or because the normal background context of the practice has changed so as to introduce uncertainty about how to proceed, or so as to introduce atypical self-interested motivations (as for example occurred with shopping practices in the Covid pandemic, resulting in individual buyers wiping out supplies and private sellers charging extortionate prices).

Decently functioning social practices, however, proceed on the presumption that participants are minimally competent social actors. They are capable of engaging in a practice *at all* precisely because they understand and comply with the constitutive rules that define the practice, and they contribute to the willing co-participation of all precisely because they understand and comply with norms of common decency. In short, decently functioning social practices proceed on the presumption that participants are compliance responsible persons. Thus, vulnerability need not be accepted nor trust invested. Rather, as Trudy Govier puts it:

One person trusts another in that she confidently expects the other to produce or respect the normal, desired events—what Garfinkel calls reproducing the “normative order of events.” Trusting in this context amounts to confidently expecting the other to act in an appropriate way. What is appropriate is defined by social custom and presumes no knowledge of the people trusted. It emerges from social experience. This is how “everybody does it,” as we have simply learned. It is how “people like us behave.” (1997, 110)

She calls this default trust “scatter trust” because it is scattered across individuals who participate in social practices, most of whom we do not know, for example,

²¹ Some philosophers have, however, explored more basic and default forms of trust. See for example Brennan 2021; Thomas 1989; Preston-Roedder 2017; Walker 2006, ch. 3; Baier 1986.

all those people involved in the food production, food safety certification, and food service industries whose competent performance we take for granted when we eat at restaurants. In adopting a stance, or attitude, of default and scatter trust, we recognize others as compliance responsible persons.

In her landmark essay, “Trust and Antitrust” (1986), Annette Baier reminds us of how pervasive reliance on others’ actual, even if possibly minimal, good will in fact is, and how very different our social life would be if we did not “inhabit a climate of trust as we inhabit an atmosphere, and notice it as we notice air, only when it becomes scarce or polluted” (234). Assuming at least minimally well-functioning social practices, we trust strangers to give correct directions, not to take advantage of us when we fall asleep on trains and planes, we trust patrons to pay for what they take from shops, trust employees to do what they are hired to do, trust people to return what they have borrowed, and trust enemies not to fire when arms are put down and a white flag is raised (239, 250).

While default trust is a prospective rather than reactive attitude, a special form of the reactive attitudes of resentment and indignation recognize compliance responsibility. Think back, once again, to Garfinkel’s breaching experiment where students behaved like boarders in their family homes. Their doing so provoked “astonishment, bewilderment, shock, anxiety, embarrassment, and anger” (1963, 226). Family members responded by saying things like “What’s gotten into you!?” and “Are you out of your mind or are you just stupid?” We are all familiar with similar expressions: “What were you thinking?” “I can’t believe you just said (did) that,” and tellingly, “Didn’t your mama teach you...?” Infractions of basic norms elicit a distinctive kind of incredulous resentment that such an infraction on the part of a competent social actor could have taken place at all. To make sense of how a presumed compliance responsible person could be doing such things, some family members treated the behavior, at least at first, as a comedy routine. Others tried to make sense of the behavior by speculating that the student was sick, or overworked, or had had a fight with their fiancé, thus volunteering excuses or exemptions (227). In short, resentments are not of a piece. Some react simply to the violation of a *normative* expectation. Others react to the violation of a *normative-cum-predictive* expectation.

There is also an important connection between the responsibility-recognizing attitudes connected with accountability responsibility—resentment, indignation, guilt, shame—and compliance responsibility. Strawsonians often take the “reactive attitudes” toward failures to meet normative expectations—resentment and indignation—to have a call-and-response structure (Macnamara 2015). Expressed

resentment and indignation communicate normative demands (or at least the fact of norm violation) and call upon the norm violator to give uptake to those demands through self-reactive attitudes of guilt and shame and to respond in appropriate ways, for example by apologizing, making amends, and undertaking character reform. The work that gets done through expressed resentment and indignation is not just the retributive work of giving people the condemnation they deserve. It is also the instrumental, agency-improving work of getting people to take seriously their norm violations and take responsibility for repairing the damage they have done and for repairing their own character.²² Central to improving agency is that the recipient of resentful and indignant calls *gets the normative message* and not just a coercive threat of social reprisal.

But, as I said in the last lecture, accountability responsible persons include those who have the realized capacity to detect and act on normatively relevant considerations, but who just don't care to do so, and those whose capacities are developable but not yet realized. That is, accountability responsible persons may be either unwilling or unable to give expressed blame the normal uptake. For those individuals, expressed resentment and indignation lack a *call-and-response* structure. At most, the unwilling can be pressured into better behavior through threat of social sanction, and the outsider to "our" social norms can be educated about what the norms are and possibly also pressured into better behavior through threat of social exclusion. But in neither case is a normative-cum-predictive expectation—of uptake in guilt or shame, or expectation of self-initiated (that is, non-coerced or pressured) efforts to apologize, make amends, and reform—appropriate.

Resentment and indignation have a real call-and-response structure, and some realistic claim to being useful in improving agency, only among those who share normative understandings and on whom those understandings typically have some motivational hold—in short, only in social worlds where individuals are both accountability and compliance responsible. Herbert Fingarette (1966) observed that "to appeal at all, we must always, finally, appeal to some acknowledged responsibility, perhaps tacitly accepted, but in any case some responsibility which the individual does accept, something for which, and in the spirit of which, he cares. Ultimately, we make this deepest and broadest appeal to the person as a responsible person—then we must wait" (73–74). Compliance responsible persons at least care about living up to familiar, basic normative expectations. They thus have a cognitive

²² See Macnamara 2013; Vargas 2013; McGeer 2019; 2017.

and motivational groundwork for having their normative competency expanded through practices of holding responsible.²³

2.4 Compliance-Responsibility-Recognizing Practices

Given that the expectation of compliance responsible persons is a normative-*cum-predictive* one, and that the compliance-responsibility-recognizing attitude is default trust that social participants will behave as they ought by living up to basic normative expectations, it should be unsurprising that the relevant responsibility practices involving not-doings. We treat people as compliance responsible persons by not subjecting them to surveillance, not checking up to see if they are doing what they are supposed to be doing, not reminding them of what they've promised to do, not doing self-protective things such as requiring formal contracts, and not avoiding them for fear of what they might do. We leave compliance responsible people alone and trust that they will manage to do, on their own steam, what is normal for competent social participants to do.

In discussing accountability responsibility practices, I indicated that there are distinctive ways of insulting people by not treating them as accountability responsible. While probably no one minds not being held to account, they might well take affront at being written off as hopelessly uneducable or hopelessly irreformable, thus not the kind of being who could come, with others' help, to be fit to participate in norm-governed practices. This is the insult that enslaved Africans received under American slavery and that criminal offenders may receive during and after incarceration. Accountability responsible persons may also take affront at being treated as though they lacked what it takes to hold others' accountable or to challenge the legitimacy of prevailing normative expectations.

Treating fellow social participants as though they were not compliance responsible persons is similarly insulting. As Jason D'Cruz observes, "Distrust is deeply dishonoring, and distrust without warrant risks insulting, demoralizing, and disempowering" (2019, 934); and "in settings in which trusting behavior is the norm and distrusting behavior is therefore conspicuous, the risk of marginalization and dishonor represents significant, even dramatic losses for the party who is wrongly distrusted and for the relationship as well" (941). Among fellow social participants, trust that others will comply with basic normative expectations *is* the norm since it amounts to trusting that others are *bona fide* social participants *at all*. The particular dishonor of not being trusted to comply with constitutive rules and norms of common decency is the dishonor of having one's status as a compliance responsible

²³. See McGeer 2015.

person denied. This is the insult that Black individuals receive when they are treated with suspicion in all sorts of social practices—driving while Black, shopping while Black, and so on. Although less insulting because not connected with prejudicial stereotypes, this is the insult one might expect telemarketers to experience when their calls are routinely monitored for “quality control,” as though they cannot be trusted to execute their basic job responsibilities. It’s the insult I experienced when Walmart instituted spot checks of shopping carts and receipts as patrons exited the building, as though patrons could not be trusted to be shopping rather than shoplifting. It’s the insult that many of us may feel at being subjected to ever more elaborate performance reports and performance assessments as though, again, we cannot be trusted to know and do our jobs.

However necessary distrustful monitoring may be or may seem to be both for the profitability of businesses and for the protection of others from the possibly untrustworthy, the expansion of distrustful monitoring threatens to erode the infrastructure that supports compliance responsibility. On the one hand, as D’Cruz suggests, it risks undermining motivation to comply (“Why should I bother if this is how I’m going to be treated?”). On the other hand, monitoring, along with public reminders of basic norms, risk suggesting that there is not in fact general compliance with a norm, and thus that it is not in fact basic.²⁴

Responsibility practices include something else besides not doing what, if done, would signal distrust: giving the benefit of the doubt when those presumed to be compliance responsible appear to be violating normative expectations (Rawls and David 2006). There must be some misunderstanding that could be resolved, or it’s just a joke, or there must be some excuse.²⁵ Here, too, there is opportunity for insult. To be accused of violating basic normative expectations is a serious matter. It amounts to being regarded as an unfit social participant and thus lacking the status “compliance responsible person.” To be immediately accused of norm violation, rather than being given the benefit of the doubt, is to be treated as someone of whom basically good behavior was not predictively expected in the first place; one was not presumed to have the status “compliance responsible person.”

²⁴ Bicchieri notes that “perhaps just sending a normative message may be interpreted as a sign that people usually do not conform to the desired behavior, encouraging transgressions” (2017, 152). See also Jules Holroyd’s (2007) discussion of empirical evidence that threats and rewards undermine compliance.

²⁵ Karen Stohr devotes a chapter in her *Minding the Gap* (2019) to what she calls “throwing the veil of philanthropy” by constructing charitable narratives of persons’ past, present, and future behavior. What I’m suggesting here is that it’s part of treating people as compliance responsible that the veil of philanthropy is initially thrown and tested before blaming people for violating basic norms.

2.5 Qualifications

Lest all this talk about compliance sound a bit too rosy, let me add some qualifications. The default presumption that participants are compliance responsible is a presumption that they will comply with *basic* normative expectations—what I’ve called constitutive norms and norms of common decency. Compliance with constitutive norms is necessary for participating in the practice at all, since constitutive norms are practice-defining (recall that someone who pockets goods in a store isn’t shopping), and willing participation in practices with others is, plausibly, dependent on being treated with common decency. So, what’s presumed is basic social competence, not compliance with more “elevated” norms, and certainly not virtue.

Further, the default presumption that fellow social participants are compliance responsible is perfectly compatible with thinking that *some* social participants aren’t compliance responsible persons at all.²⁶ It is also compatible with recognizing that some social participants will not be compliance responsible with respect to particular practices: they may be cultural outsiders who have not been trained into “our” norms; or they may be not-yet-socialized novices to a practice (for example, the medical student not yet socialized into the norm-structured practice of medicine); or they may simply be outsiders to a practice (for example, the non-Catholic attendant at a Catholic mass).

In addition, the default presumption is not that compliance responsible persons are flawlessly compliant, just for the most part so. Given that a presumption of compliance responsibility doesn’t entail that everyone always does comply with constitutive norms and norms of common decency, taking precautionary measures may be prudent, especially when the stakes are high. So, shopkeepers install security cameras, airports search bags and persons, tourists wear money-belts, women avoid walking alone at night on city streets, Black people avoid police interaction, internet sellers request CVV numbers, potential employers conduct background checks, and transactions may proceed only on the basis of enforceable contracts.

It’s also important to keep in mind that the range of norms that count as basic varies from practice to practice. The *basic* normative expectations within friendships or monastic life may differ substantially from the basic normative expectations for using public transportation or engaging in professional politics. What is a mere matter of common decency among friends may well not be among strangers. I haven’t meant to suggest that there’s some single standard for common decency that applies universally

²⁶ Something similar is true for the default presumption that people are accountability responsible. That presumption is compatible with some, for example, psychopaths, turning out not to be.

to all social participants (although there are surely *some* norms that do apply nearly universally, such as prohibitions against murder). The claim has only been that if a practice is decently (i.e., non-coercively) functioning, the default presumption must be that practice participants are compliance responsible with respect to whatever constitutive norms and norms of common decency apply to that practice.

I have also not meant to suggest that basic practice norms are necessarily ones that, from a critically reflective, non-participant stance, we would regard as *legitimate*. The norms structuring practices, especially in hierarchically structured societies, can be defective. The constitutive norms and norms of common decency may concern only “insiders” or “privileged groups.” So, for example, the scope of the constitutive norm of policing that police not murder citizens may not extend to some social groups. The Black Lives Matter movement responded to this kind of defect in U.S. policing practice. Similarly, the #MeToo movement responded to the failure to include a prohibition against sexual assault among the constitutive or common decency norms in many practices. What this means is that basic trust in *compliance* is not the same thing as trusting practitioners not to injure or wrong those who fall outside of the scope of its basic practice norms.

Finally, I haven’t meant to suggest that *all* of our social practices operate on the presumption that participants are compliance responsible. It is possible, though very costly, to keep a social practice afloat among non-compliance responsible people. Institute enough surveillance mechanisms and punitive sanctions and compliance can be coerced. But the larger the number of social practices in which no default presumption of compliance responsibility operates, the more one might be inclined to say that social life has broken down. In short, it would not be possible to have anything resembling the social life characteristic of decently functioning societies absent a default presumption that people are compliance responsible.


2.6 Conclusion

I began by pointing out that a number of features that are often thought important to being a responsible person are not features of being accountability responsible. Those include a presumption that fellow participants in social practices have a basic normative literacy, such that many of our normative expectations about how people ought to behave are also predictive expectations about how typical social participants will behave. Given basic normative literacy, the attitudes and practices by which we hold people responsible have a call-and-response structure. We call on people to live up to normative expectations with which they are familiar and expect that such calls will generally receive uptake. Also, given basic normative literacy, it

will generally be true that those who disappoint our normative expectations will be liable to blame.

Because accountability responsibility is a “ground floor,” non-social conception of responsible persons, we need to add a second dimension of responsibility to capture the way we think about responsible persons within shared norm-structured social practices. Socialization and social experience produce social participants who are compliance responsible in the sense that they are familiar with, and largely disposed to comply with, basic social norms. In our everyday life together, we simply assume that fellow social participants are compliance responsible. “Compliance responsible person” is thus a *default* status assigned to fellow social participants.

Focusing on compliance responsibility brought into view a distinctive set of responsibility-recognizing attitudes and responsibility-recognizing practices, and with those, distinctive ways of insulting people by not recognizing their status as compliance responsible persons.



3

TAKING RESPONSIBILITY

“It seems worrisome to account for morally responsible agency in terms of only some of the ways we engage people (including ourselves) as responsible agents, in terms of only some of the attitudes that characterize our responsibility practices.” *David Beglin (2020, 2349)*

The guiding idea of these lectures is that if we want to understand responsibility, including the full range of our responsibility-recognizing practices and attitudes and the distinctive ways that responsible persons can be insulted, we should begin by examining what it means to have the *status* “responsible person” rather than by examining more specifically what it means to be responsible *for* a particular action or consequence. Thinking about responsibility *for* naturally invites us to focus on responsible persons’ vulnerability to blame and general answerability for deeds that appear to violate normative expectations. That in turn invites us to focus on those responsible persons who share a normative community with us and thus who might sensibly be called upon to render an account that we can understand, and might sensibly be blamed for normative failures that they can understand.

Thinking about responsibility *for* also naturally invites us to think about responsible persons primarily as people who can *fail* to meet our expectations, rather than as people who generally *meet* our basic expectations and can be trusted to do so.

The result is that we miss out on important ways that we interact with those who are not yet part of our normative community as nevertheless responsible persons. We also miss out on the most commonplace ways that we interact with those who are part of our normative community and who, having been socialized into it, generally comply with its basic norms. A related result is that the negative reactive attitudes of resentment and indignation come to occupy an unduly prominent place in accounts of responsibility, as do blaming and sanctioning treatments. A central aim of the

first two lectures was to draw attention to responsibility-recognizing practices and attitudes that are in fact central to interacting with others as responsible persons, even if they are not perspicuous from the point of view of thinking about persons' responsibility for normative derelictions.

Although significantly departing from Strawsonian approaches to responsibility—with their stress on answerability, liability to blame, normative demands, and so on—I have meant to leave in place Strawson's central insight, namely, that we are invested in responsibility because we are social beings whose lives are oriented around participating with others in norm-structured practices. We thus have a basic concern that we be able to share social practices with others, and thus a basic concern with the capacities that, actually or potentially, fit the persons around us to be social fellows.

A second dimension of the Strawsonian approach that I've meant to leave in place is methodological. If we want to understand responsibility, we need to begin by reflecting on what social life with others is like, rather than beginning from metaphysical views about contra-causal freedom or determinism. Taking up this idea, I've suggested that we should attend to the *default* status of being a responsible person that our social interactions with others presuppose. Those social interactions will typically take place within ongoing social practices, within a shared normative community—but not always, since there will also be outsiders both to our larger normative community and to practice-specific communities with whom we also engage as responsible persons.

So far, I have not questioned the standard assumption that being a responsible person is entirely a matter of having the capacities that fit one to stand under *obligation*—that is, to be subject to normative expectations. I have only made adjustments to that view, emphasizing that the relevant capacities may be merely developable, and emphasizing the automaticity of much norm-compliant behavior within ongoing social practices. Those adjustments meant substantially enlarging the range of responsibility-recognizing practices and attitudes. Nevertheless, the core underlying thought that responsibility and obligation go hand in glove remains the same. The idea that to be a responsible person is to be liable to blame for failure to do what one is obliged to do is also largely preserved. Even if we hopefully attempt to scaffold not-yet-(fully)-developed reasons-responsiveness, resentment is waiting in the wings once the scaffolding is complete, ready to emerge when norm violations occur. Even if we trustingly taking compliance with basic norms for granted, resentment is waiting in the wings ready to emerge should norm violations occur.

There is however this potential difficulty with tying responsibility so tightly to obligation: people often elect to do good things that are normatively optional. Indeed, as I'll elaborate shortly, the social world we are familiar with is pervasively structured around such elections. There are also responsibility-recognizing practices and attitudes associated with good-doings. To mention only a few: asking for volunteers, sending thank you notes, and feeling and expressing appreciation and gratitude.

3.1 Accommodating Positive Reactive Attitudes

A number of philosophers, including myself, have already raised doubts that what Strawsonians call “positive reactive attitudes”—appreciation, gratitude, praise, admiration, approbation—fit models of reactive attitudes developed specifically to capture distinctive features of blaming attitudes like resentment and indignation, whose expressions hold people to normative expectations that have been disappointed.²⁷ A conception of responsible persons designed to explain what licenses resentment and indignation may not be the conception we need to explain the appropriateness of appreciation and gratitude.

Initially it might seem as though resentment and appreciation (to pick focal negative and positive attitudes) are entirely symmetrical. Resentment recognizes responsibility for *failures* to meet normative expectations. Appreciation recognizes responsibility for *exceeding* normative expectations. So, both refer to normative expectation. But I don't think much hay can be made of this fact.

“Obligation,” when used synonymously with our everyday normative-cum-predictive notion of normative expectation, is a social concept. It is the concept of what one person may properly demand of others within a shared normative community, where a community might be an entire society, a neighborhood, members of an occupation, and so on. Similarly, the normatively optional—the supererogatory—is a social concept. It is the concept of what we may not properly demand of others despite its being a good thing to do (Ferry 2015; Levy 2015). Actual social worlds will vary in what they take to be demandable versus optional, and how much of persons' behavior is subject to demands versus how much is left to voluntary election. In principle, a social world might regard nothing as optional other than selection among equally weighted obligations. By our lights, this would be an overly demanding world. In principle, a social world might regard nothing as demandable. Perhaps a community of people who are especially good judges of what would be good to do, and who are highly motivated to do it, might see no point in making normative demands on

²⁷ See Telech 2020; Eschleman 2015; Calhoun 2022; and Macnamara 2013; 2015.

each other. These observations suggest two things. First, the normatively optional does not make *essential* reference to obligation. That is, it is not necessarily a notion of exceeding obligation. The community of highly motivated good judges might appreciate and feel grateful for each other's various elections of good things to do without conceiving of these good deeds as exceeding normative expectations, since they do not operate with the notion of normative expectation. Second, ordinary social worlds take the range of good things that might be done and draw a line: on one side are the *non-optional*; on the other are the *non-obligatory*. Whenever societies engage in such line drawing, *both* sides of the line make some reference to their opposite side. So, not much hay can be made of the fact that the normatively optional (in the sense of supererogatory) *exceeds normative expectations*.

In short, the fact that, in typical societies, the normatively optional exceeds normative expectations is not *by itself* a reason to think that resentment and appreciation share a similar structure. It might nevertheless be true that they do. If so, they would embody the same conception of "responsible person." That means that whatever capacities we conclude responsible persons must have to be subject to resentment will be the very same capacities responsible persons must have to be subject to appreciation. There would then be no point in paying special attention to the positive responsibility-recognizing attitudes, since they will tell us nothing about the nature of being a responsible person that we hadn't already learned by attending to negative attitudes.

But does appreciation simply mirror the structure of resentment? Here is a simplified, but I think recognizably Strawsonian, account of resentment: Resentment (1) is a response to others' failure to live up to normative expectations; resentment (2) incipiently communicates demands concerning what was originally expected but not received, as well as demands for appropriate responsive behavior to this failure;²⁸ and expressed resentment (3) comes with an "implicit RSVP" for uptake of these demands in such self-reactive attitudes as guilt, remorse, and shame, as well as in reparative actions.²⁹ Resentment, thus described, has a complex communicative structure. It seems evident enough why it should. Failures to meet normative expectations are serious matters. Such failures treat persons without the regard or respect owed to them, and thus are intrinsically insulting. They also threaten the smooth functioning of social practices on which we, as social beings, depend. It's thus important to underscore for the norm violator exactly what was normatively

^{28.} As Coleen Macnamara (2015) has argued, even if resentment isn't always expressed, the attitude is "for communicating" in much the same way that an unsent email is for communicating, or a "no trespassing sign" in the local hardware is for communicating. This is what I mean by "incipiently communicates."

^{29.} The phrase "implicit RSVP" is Darwall's (2006).

expected, and possibly also why it was. Once attention has been brought to the fact of norm violation, additional demands can be explicitly pressed or implicitly made—demands to acknowledge fault, apologize, make amends, engage in whatever reform is necessary so that the violation doesn't happen again, and so on. The point of making such demands is, of course, to get the norm violator to respond appropriately, hence the importance of communicative uptake in self-reactive attitudes and subsequent reparative activity.

Obviously, appreciation and gratitude don't have this exact same structure since they are responses to good action that is not normatively expectable, and thus demands won't figure into those attitudes. But they might have the same *general* communicative structure of (1) being a response to normative performance, (2) being incipiently communicative, and (3) coming with an RSVP for uptake of that communication.

Both Coleen Macnamara and Daniel Telech spell out how such an account of, respectively, gratitude and praise might go. Macnamara (2015; 2013) argues that gratitude *recognizes* the positive moral import of what was done, and that expressions of gratitude are recognitive speech acts. Gratitude seeks uptake of that recognition in the target's coming to see herself as she is seen and thus to feel self-approbation. Self-approbation motivates communicating receipt of the message by, for example, saying "You're welcome." Self-approbation also motivates doing more good things that build relationships and bind the moral community together (analogously to the reparative activities motivated by guilt). Telech (2021) proposes something similar. Praise communicates the significance for the praiser of what was done, and issues an *invitation*. The invitation is to accept credit from the praiser and join the praiser in *jointly valuing* the praisee's act. The praisee gives emotional uptake to that invitation by feeling self-directed pride and by such discursive communicative responses as "You're welcome!" or "I'm happy to have helped" (analogous to apologies).

While we learn something about the *varieties* of moral address and responses to that address by attending to gratitude and praise (in particular, moral address includes recognitives and invitations), we don't learn anything about the nature of responsible persons that we hadn't already learned from an examination of resentment. For Macnamara and Telech, both the positive and the negative reactive attitudes presuppose that responsible persons have the requisite capacities for engaging in the discursive and behavioral exchanges characteristic of *accountability* responsible agents. Expressed resentment holds accountable agents *to account* for misdeeds by communicating that they need to acknowledge their misdeeds (in feelings of

guilt, shame, remorse) and to do something in response to that acknowledgment (apologize, repair the harm). Analogously, expressed gratitude holds accountable agents *to account* for good deeds by communicating that they need to acknowledge their good deeds (in feelings of pride, self-approbation) and to do something in response to that acknowledgment (say “You’re welcome,” do community building, join the praiser in valuing the deed).

I don’t think this is right. Pressing gratitude and praise into the mold of accountability holding moral-address-and-responsive-uptake produces strained accounts of these attitudes. No doubt some expressions of gratitude and praise aim to get the agent to appreciate what a good thing she’s done and to feel pride and self-approbation. Sometimes the person who has done a good thing doesn’t realize quite how good a thing they’ve done. In that case, praise might indeed invite the praisee to upgrade their assessment, thereby coming to see themselves as they are seen and to jointly value what they have done. Such upgraded assessment might well result in intensified pride. But this is an odd view of what gratitude and praise in general are like. To see this, note that resentment aims to get the norm violator to see themselves as they are seen, because norm violation is motivated by such things as inattention, self-interest, and sometimes plain indifference to others’ interests. Thus, norm violators simply aren’t paying sufficient attention to the normatively relevant features of their actions. Moral address calls those features to attention. By contrast, those who exceed normative expectations typically elect to do so precisely because doing more seems a good thing. They are not inattentive to the normatively relevant features of what they do, and they do not need the goodness of their actions brought to their attention through others’ recognitives or invitations to jointly value. Indeed, they deserve gratitude or praise only if their election springs from some sensitivity to the goodness of what they do. We can agree that expressions of gratitude and praise do *something*. But what they don’t generally do is focus the target’s attention on something they weren’t already aware of in the first place.

So, construing expressions of gratitude as attention-focusing moral address seems strained. So too does construing them as aiming for *emotional* uptake. Having elected to do good, the agent already has grounds for pride and self-approbation. So, the agent’s emotional “uptake” typically precedes communicative exchanges; it is not reasonably an aimed-for effect of them. Indeed, I suspect it’s more common for people to complain that their good deeds, of whose goodness they are well aware, haven’t received proper uptake via *others’* emotional responses of gratitude and appreciation.

The idea that gratitude and praise aim for *behavioral* uptake seems equally strained. Etiquette norms may require or advise acknowledging thanks with “You’re welcome,” or minimizing the act with “Think nothing of it” so as not to burden others with a duty of grateful repayment. It’s hard to see how either discursive response centrally concerns acknowledging that the person has got the message that she’s done something good. “You’re welcome” and “Think nothing of it” simply aren’t analogous to “I’m sorry”; the latter *does* convey “message received.” In addition, aiming to get uptake of grateful messages in *more* good deeds, motivated by self-approbation, seems distinctly *ungrateful* (and possibly manipulative).

If positive reactive attitudes don’t easily fit the communicative model designed for negative reactive attitudes that recognize accountability, perhaps they don’t embody the same conception—or as I would say, dimension—of being a responsible person. This is what I propose: just as basic trust recognizes others as *compliance responsible* persons, gratitude, appreciation, and the like recognize actions that befit *responsibility taking* kinds of beings. I turn now to the third dimension of being a responsible person—responsibility taking.

3.2 Responsibility Taking and Social Practices

One might be skeptical of the idea that being a responsibility taker is part of the default status “responsible person.” To take responsibility—as I am using this phrase—is to elect to do good things that are not normatively expectable. One might think such elections do not depend just on being a responsible person, but on some added extra that some responsible persons bring to the table. They are virtuous in some way. Perhaps they are especially concerned with others’ welfare. Perhaps they are especially conscientious about doing their duty and so regularly end up doing more than required—they have the virtue of being highly responsible persons. Or perhaps they personally value particular ends, like the success of an academic program or a political campaign, and so are willing to take on responsibility for promoting those ends and to execute those assumed responsibilities in especially diligent, time-consuming, and creative ways. They are, as we might say, personally committed. In addition, talk about the *supererogatory* invites reflection on rare paragons of supererogatory behavior—saints and heroes—and to suppose that those who exceed normative expectations in less flashy ways (like doing favors) similarly have a something-extra in them besides just being responsible persons. And so, once again, we might think that, if we want to understand what it means to have the default status “responsible person,” we should focus on the baseline capacities for standing under normative expectations—capacities connected with resentment and basic trust.

But now consider this fact about social life as we know it: in everyday social life together, we take for granted not only that social participants have what it takes to participate in norm-structured social practices, but also that they are disposed to, and sometimes will, elect to promote the goods that those practices serve in ways that aren't required by existing norms. That presumption is *pervasively* built into our social practices. Member organizations like churches, clubs, and public radio corporations rely, in part, on the users of their services voluntarily contributing money or labor. Charitable organizations serving the needs of some group—such as food pantries, homeless shelters, animal rescue organizations, and international humanitarian aid organizations—similarly rely on donations of money and labor. Public-good-providing institutions such as public libraries, museums, and national parks rely on volunteers to give talks and tours. Formalized groups that engage in political advocacy, such as campaign organizations, the National Rifle Association, and the Humane Society Legislative Fund, rely on economic contributions and volunteer labor as well. All these groups take for granted that *appealing* for funds and volunteers is a productive thing to do.

Not only do many organizations rely on voluntary contributions of money and labor to support their work, but many places of employment distribute tasks among their employees by asking for volunteers rather than mandating that specific employees add these tasks to their existing job duties. Think for example of how common it is in academia to ask for volunteers to serve on various committees, or to oversee some new initiative. Such requests may be general calls for volunteers or directed appeals of the form “Would you be willing to be the one who...?”

In the domain of academia, some functions essential to academic life get done entirely by making requests that no one is obligated to accept. Among those essential functions are manuscript reviewing and serving as an external referee for tenure and promotion cases, and, in many countries, giving uncompensated lectures at other institutions.

Perhaps one of the most important assumptions made across social life is that those charged with the responsibility to look after something will typically not be so stingy with their time and energy as to do merely the bare minimum necessary to discharge their responsibilities. Parents have the responsibility to look after their children, graduate directors the responsibility to look after graduate programs, pet owners the responsibility to look after their pets, and employees to look after what falls within their purview, even when it is not part of their formal job duties. What counts as satisfying the *demandable* normative expectations for satisfactory execution of these responsibilities is much lower than the performance we often assume most will deliver.

Finally, there's the realm of everyday interpersonal interaction between both strangers and intimates where we feel free to ask for favors, ask for directions, ask for help carrying things, ask for advice, and, in the case of the impoverished, ask for money on the streets. We simply do a lot of *asking* of other people to do things they aren't under obligation to do.

What is striking about all this is that we simply assume that requesting, appealing, and relying on others to do more than the required minimum makes sense. So, we arrange our norm-structured practices in such a way as to leave many of the things that make for successful social life together up to individual election. In short, the social practices we are familiar with aren't just norm-structured; they are norm-*and-election* structured. Social practices need not take this form. They might have been designed for a society of beings who are merely fit for living up to normative expectations and nothing more. Successful social life together would then be achieved by vastly expanding the domain of the normatively expected so that virtually all of what needs doing is either a matter of standing obligation or special assignment. And, where we currently recognize imperfect duties to promote valuable ends on some occasions and in some ways, we would instead perfect those imperfect duties by converting them into normative expectations.³⁰

It's sometimes said that the reason for *not* arranging society this way is that it's important to preserve a realm of free action in which individuals make autonomous choices, free from the influence of threat of sanctions. Michael Ferry (2015), for example, observes that, "if every moral ought were properly enforceable by means of demands and punitive moral sanctions, very few of our significant decisions would be off-limits, and our range of truly free choice would be severely restricted" (63). But this isn't a convincing reason for having partially election-structured social practices, since, absent a responsibility-taking population, much of what needs to get done wouldn't get done via free election. To my mind, the fundamental reason for creating a realm of the normatively optional is that not doing so would insultingly treat responsible persons as though they lacked an important dimension of being a responsible person.

In sum, we have a variety of reasons for thinking that there is a third dimension to the default status "responsible person" that concerns being able, and sometimes disposed, to take responsibility by electing to promote good ends in ways that

³⁰ Alan Buchanan (1996) argues that there are good reasons for sometimes perfecting imperfect duties, such as the way the indeterminateness of what is to be done encourages moral laxity, the reluctance to promote good ends when there's a lack of assurance that others will do likewise, and the inefficiencies produced by leaving what is to be done up to individual judgment.

aren't normatively expected. First, there are the positive responsibility-recognizing attitudes of appreciation, gratitude, praise, admiration, and approbation. These are ordinary, commonplace attitudes. But they do not fit the communicative model of negative reactive attitudes that hold miscreants accountable. That is some reason to think they embody a different conception of responsible persons than those connected with accountability. Second, social life as we know it is pervasively structured around the presumption that social participants are responsibility takers. This includes socially drawing a line between the normatively expected and the normatively optional. Third, and relatedly, social practices make use of a set of *prospective* responsibility-recognizing practices: asking for volunteers, making appeals, issuing directed invitations and requests, asking for favors or help, relying on those who have responsibilities to look after some good to do more than the bare minimum. Those prospective responsibility-recognizing practices are connected with a prospective responsibility-recognizing attitude: hope. It is a hope that when we call on them to volunteer, to donate, to do favors, to join social movements, they will have adopted or be willing to adopt a particular good as something they are willing to take responsibility for promoting. It is a hope that in fulfilling their assigned responsibilities, they will value the goods served by those responsibilities enough to do more than is minimally required.

Practices that recognize responsibility having been taken are also very much part of social life as we know it. These include not only the interpersonal ones of saying "Thank you," "I appreciate it," "What a lovely thing to do," and the like. They also include a wide array of institutionalized recognitions. The U.S. recognizes an enormous number of days, weeks, and months dedicated to appreciating the good that individuals electively do, including exceeding the normative expectations of particular jobs and roles: Teacher Appreciation Week, Martin Luther King, Jr. Day, Veterans Day, Customer Service Week, National Non-Profit Day, National Volunteer Week, Employee Appreciation Day, Mother's and Father's Day. Journals publish "In Appreciation" essays for former journal editors. Businesses bestow service awards and merit raises on employees for doing more than required. Memorial eulogies mention the deceased's positive contributions. These are all institutionalized occasions for expressing the responsibility-recognizing attitudes of appreciation, gratitude, praise, approbation, and admiration.

3.3 Responsibility Taking

Perhaps the strongest reason for thinking that responsibility taking is part of the default status of being a responsible person is simply the oddity of there being responsible persons fit to participate in norm-structured practices who don't also

have the capacity and disposition to promote goods in non-required ways. Social practices arise because there are things we collectively value. Consider the huge array of social practices oriented around the goods of family life, health, play, safety, transmission of knowledge, social inclusion, and access to income and other material resources. The good of play, for example, underwrites such practices as puzzle working, playing games, practicing sports, producing and marketing toys, holding professional and amateur competitions, employing and being employed as a trainer, and establishing recreational centers and athletic camps.

Practice norms specify what people are to do in order to realize or promote a particular good. For example, the point of both norms governing hospital practices (such as doctors' and nurses' job responsibilities and medical codes of ethics) and the etiquette norm of covering your mouth when you sneeze is to secure the good of health. If we assume that social participants are able and disposed to conform to the norms of social practices, we surely equally assume they can understand the goods served by those practices. That is, they can understand the *point* of being normatively expected to behave in particular ways. If they grasp the point of what they are required to do, they have what it takes to see that they could also *take* responsibility for promoting that point in ways that aren't required by practice norms. It is thus difficult to imagine social participants who are fit to participate in the *required* aspects of norm-structured practices who aren't also equipped to promote the practice's underlying good in *non-required* ways.

Perhaps the most familiar ways of doing so are to do more things than your job or role requires and to execute what you are required to do in more diligent, time-consuming, creative, or taxing ways than is minimally required. Teachers are required to create educational classes. Some teachers elect to take webinars on improving teaching and choose to implement the latest pedagogical techniques even though taking such webinars and employing the latest pedagogical techniques aren't required for the job. In doing so, they take responsibility for seeing to it that students are well educated. One can also take responsibility for promoting a good by *staying* in a burdensome job that one could quit. During the Covid pandemic, many health care workers elected to stay in jobs that had become highly risky and involved longer hours rather than quit. They took responsibility for seeing to it that the sick were cared for.

Another way of taking responsibility is to volunteer to, or to accept a direct request to, take on a new responsibility, agreeing, for example, to see to it that the new pet gets walked or that specific students are mentored, or promising to help a friend

move. Responsibility can also be taken without making agreements or promises. The essayist David Sedaris, for example, became even more famous when he elected to spend three to eight hours a day picking up litter. He didn't promise to do this or continue doing it. Activists in the BLM movement took responsibility for challenging police practices; and activists in the #MeToo movement took responsibility for publicizing and calling out sexual harassment. More simply, one might take responsibility for others' welfare by stopping to help obviously lost tourists. In all of these cases, there is something that needs doing in order to promote a good. Responsibility takers answer the question "Who is to do what *someone* needs to do?" by affirming "I will!" They self-assign responsibilities.

Taking responsibility is important to our social life together because existing practices often do an imperfect job at securing the goods that explain why we have those practices in the first place. That imperfect job may be a result of the fact that we don't have a sufficiently expansive set of practices. Think, for example, of the various practices in a graduate program aimed at student success—classes, supervised independent study, mentoring, reading groups, job placement services. But a particular graduate program may not have a practice of offering workshops on how to publish. Expanding the range of practices aimed at student success depends on someone taking responsibility for changing how things are done.

Practices may also do an imperfect job at securing their underlying good because they lack, or fail to properly enforce, norms that are important to securing the good. Police practice, for example, may not preclude potentially lethal behavior, or may fail to properly enforce existing norms precluding lethal police behavior. The good of citizen safety thus fails to be adequately secured. In such cases, individuals (for example, whistle blowers) and groups (for example, the Black Lives Matter movement) may take responsibility for effecting changes to practice norms.

Even the best practices will not make all the possible ways of promoting its good a matter of obligation. Norms directing people as to what they are required to do simply don't cover all that it is good to do.

Obviously, it is not part of our *default* conception of social participants as responsible persons that they *will* take responsibility for promoting specific goods,³¹ and certainly not that they will take responsibility on *all* occasions where they

³¹. This is a generalization to which there can be local, practice-specific exceptions. Leaders of an organization may be selected precisely because they are responsibility takers, and service organizations may be largely governed by norms of responsibility-taking. A more demanding default conception may thus apply to leaders and members of practices devoted to service.

might. Nevertheless, our social life together, at least as we know it, proceeds on the *appellative*³²-*cum-predictive expectation* that social participants are capable of electing to promote the good and will elect to do so on at least some occasions. They have the capacity to appreciate the goods underlying particular practices and to adopt those goods as their own ends; they also have the cognitive capacity to see how a good might be promoted in ways that are not required, and thus to use their own discretionary judgment.

Although having responsibilities is a burden, being fully recognized as a responsible person means being granted the option of taking on responsibilities. One way of insulting responsible persons is to omit them from requests to volunteer or to persistently rebuff their offers to help when others' offers are accepted. This is the insult sometimes suffered by the elderly. There is also the insult of pressuring responsible persons to electively do good by misrepresenting the elective deed as normatively expected. For example, Alley Cat Allies, a cat rescue organization, began issuing pink "overdue" notices to its supporters who had failed to "renew their memberships." Similarly, non-profit board members are often asked to make meaningful voluntary contributions that are simultaneously described as expected of all board members. More commonly, responsible persons experience the insult of not having their responsibility taking publicly acknowledged and appreciated. Their elective efforts to do good things that they weren't obligated to do get ignored, dismissed as their simply satisfying personal desires or needs, or misconstrued as merely fulfilling normative expectations.


That it matters both to be recognized by and to recognize others as responsibility taking kinds of beings helps to explain why people so often offer what appear to be excuses or justifications for declining invitations and appeals to do what is in fact normatively optional. "No"—and worse yet, "No, I don't want to"—seem inappropriate responses to invitations, requests, and appeals. But if the invitation, request, or appeal is to do the normatively optional, why aren't these appropriate? Why offer instead "I'm sorry, I can't," "I already gave," and the like? One possibility is that people often mistake the normatively optional for the normatively required, and so aim to deflect blame with excuses and justifications. But then we must suppose that such confusion is widespread. Another possibility is that people are being insincere in offering excuses or justifications, knowing full well they need not do so. But then why do others, who also know full well that the requested behavior

³² I take this term from Michelle Mason (2017), who distinguishes between normative demands and aretaic "appeals to comport oneself in the manner befitting" some ideal, e.g. of various special relationships (156). She has something a bit different in mind than I do, since her aim is to capture the resentments we feel toward those who "violate" ideals (167).

is optional, play along with insincere excuses and justifications? More plausibly, Gregory Trianowsky (1986) suggests that offering excuses reflects a concern with appearing to fall short of what the fully virtuous person would do. As Michael Ferry (2015) suggests, omitting the supererogatory “may still be seen as a kind of moral failure from the perspective of those reasons that directly recommend the act” (62); in which case, excuses and justifications might deflect criticism for failing to do what one ought, even if what one ought to do is not required. There might, however, be something simpler and more basic at stake that motivates the offering and acceptance of excuses and justification where none are necessary. The offering sends the message: “I understand the good at stake, and my declining shouldn’t be taken as evidence that I’m not a responsibility-taking kind of being who is unable to appreciate the good and elect to promote it.” Accepting the unnecessary excuse or justification returns a message of acknowledgment of one’s status as a responsibility-taking kind of being. It’s not one’s *virtue* that needs defending, but one’s default status as a responsible person.

3.4 Conclusion

I’ve argued that our social practices are pervasively structured on the presumption that social participants have the capacity and disposition to elect to promote the good that underwrites those practices in non-required ways. I’ve drawn attention to practices and attitudes that recognize others, both prospectively and retrospectively, as responsibility-taking kinds of beings. The aim was to make plausible the claim that the *default* status “responsible person” includes being a responsibility taker. While it is true that those who take responsibility sometimes display noteworthy virtues—of generosity, compassion, conscientiousness, and so on—I have not argued that the default presumption is that social participants are virtuous agents. Rather, I have argued that the default presumption is that social participants are sufficiently equipped for it to make sense to organize social life so that only some promotion of the good is normatively expected and much left normatively optional: they are sufficiently capable of appreciating the goods served by norm-structured practices, have sufficient motivation to adopt some of those goods as personal ends, and have sufficient cognitive capacity to see how those goods might be promoted in non-required ways. Issuing invitations, requests, and appeals, and expressing gratitude and appreciation are, I’ve suggested, responsibility-recognizing practices. The so-called positive reactive attitudes are not, I have argued, just the positive counterpart to negative attitudes like resentment and indignation that recognize accountability responsibility. Instead, they acknowledge a distinctive dimension of being a responsible person, namely being a responsibility taker.



Summary Conclusion

In these three lectures, I've adopted a quasi-Strawsonian methodology by suggesting that we start from the facts as we know them about what it is like to engage with others as fellow social participants. I have, however, rejected the strategy of focusing narrowly on what it is like to hold others accountable within a shared normative community. That focus is too narrow because some of our engaged interactions are with others who are not party to our shared normative community. It is also too narrow because our responsibility-recognizing practices and attitudes considerably exceed those involved in holding others responsible for failures to live up to normative expectations. As part of shifting attention from responsible persons' liability to blame, I've suggested that we focus on "responsible person" as a valuable status that individuals might *want* to have and feel insulted by not being treated as having.

Because some of our responsibility-recognizing practices and attitudes are not cognitively salient, I've recommended beginning from the question "What default conception of 'responsible person' do our social practices presuppose?" I've proposed that the default conception of responsible person presupposed by our norm-structured social practices is three-dimensional. Responsible persons are accountability responsible, compliance responsible, and responsibility takers.

Because these are three distinct dimensions, with distinct sets of responsibility-recognizing practices and attitudes, one might wonder what unifies my proposed conception of responsible persons. This unification question is not unique to my proposal. Strawsonians in general confront the unification question whenever they try to explain why the long list of reactive attitudes originally proposed by Strawson all count as reactive attitudes towards responsible persons. Besides resentment and indignation, Strawson mentions, gratitude, forgiveness, love, hurt feelings, and in general all those interpersonal reactions to not only others' goodwill and its absence, but to others' *affection* and *esteem* and their absence (2008, 5).

If we are to derive the conception of responsible persons from our ordinary interpersonal attitudes, we certainly could begin from the full list of attitudes reacting to others' goodwill, affection, and esteem. But this is not at all a promising method for getting at a conception of responsible persons. On the one hand, as some have argued (Wallace 1994; Shoemaker 2019; Beglin 2020), there doesn't seem to be anything that *unifies* the variety of interpersonal responses to others' goodwill, affection, and esteem from which we could derive a conception of responsibility. What might reciprocated affection or hurt feelings at overhearing an unflattering remark never intended for our ears have in common with resentment and indignation? Worse, why should the heterogenous collection of emotional responses arising from *interpersonal* life with others all be taken as responses to others' status as *responsible* persons? In the face of this, the options appear to be either to rely only on those attitudes—and their associated practices—that obviously concern *holding* people responsible—resentment, indignation, and guilt—thereby securing not only a unified account of what makes an attitude a reactive attitude, but also a unified account of the capacities that make a person a responsible person.³³ Or we rely on the full range of interpersonal responses and simply aggregate all the different agential capacities implicated in different reactive attitudes, giving up on the idea of a unified account.³⁴ Or we conclude that “responsibility ... is too fractured a notion to merit unified investigation,” and indeed that some of the fractured parts—such as whatever we would say about hurt feelings as a response to responsible persons—bears virtually no resemblance to any familiar notion of responsibility (Shoemaker 2019, 147). As I've argued in this lecture, even if we limit ourselves to the negative reactive attitudes of resentment and indignation and the positive reactive attitudes of appreciation, gratitude, praise, approbation, and admiration, I don't think it's possible to provide a *single* account of the conception of responsible person presupposed by both positive and negative reactive attitudes. In short, I think we should concede that our conception of responsible persons is complex and multidimensional.

Even if the unification question is not unique to my account, it's still a challenge: Why think the three dimensions are dimensions of a single thing—being a responsible person? One option, which I think is a good one, is to identify an underlying concern that unifies the three dimensions.³⁵ Strawson himself suggested that the

^{33.} Wallace proposes restricting reactive attitudes in this way in *Responsibility and the Moral Sentiments* (1994).

^{34.} Beglin mentions the aggregative approach as an option for Strawsonians who derive an account of responsible persons from reflecting on the reactive attitudes, although he thinks it's a bad option (2020, 10).

^{35.} Beglin (2020) argues for the importance of identifying the underlying concern. I have not, however, adopted his tentative proposal for what that concern is.

underlying concern is with others' goodwill and ill will or indifference. Certainly, it's possible to see our concerns with others' capacities to live up to normative expectations, to be compliance responsible, and to take responsibility as reflecting an underlying concern with the good will (and absence of ill will or indifference) of those with whom we engage as participants in an actually or potentially shared normative community.

But I think we can do better than this vague reference to good will and ill will. The underlying concern is with the possibility of sharing social practices with others. It is a feature of shared social practices that (a) they are norm-structured—thus we are concerned with capacities that enable individuals to be governed by those norms; (b) they are decently functioning only if some degree of automatic and routine compliance is present—thus we are concerned with participants' normative literacy; and (c) the norms themselves are justified by the goods they promote—thus we are concerned with their capacity to promote those goods in non-required ways. In short, we are concerned with identifying and properly interacting with those who can comply with norms, who are disposed to comply at least with basic norms, and who can advance the goods underwriting social practices in ways that aren't already specified. Given this complex concern, we should expect complexity in a conception of responsible persons and in the responsibility-recognizing practices and attitudes. Indeed, the account I've offered makes salient a much wider range of responsibility-recognizing practices and attitudes, and thus captures more of the richness of what Strawson called "the participant attitude." The account does so without shoehorning all responsibility-recognizing practices and attitudes into a single model of what a responsibility-recognizing practice or attitude must be like. Some, for example, have a communicative structure, others do not.

In addition to wondering what unifies the three dimensions of the default status "responsible person" that I've suggested, one might wonder whether the three dimensions could come apart. Might one be only partially a responsible person? I think the answer to this question is "yes."

First, one might be merely accountability responsible and neither compliance responsible nor a responsibility taker. This is because being accountability responsible does not depend on being part of a shared normative community. One might thus be accountability responsible but be unfamiliar with both the norms structuring social practices and the goods they serve. If you are an outsider to a normative community—whether that be a whole society or a particular social practice, you will not share the relevant normative understandings and will not

have been socialized into them. Thus, it is possible to be accountability responsible without being compliance responsible within a particular society or a particular practice. Being unfamiliar with the relevant society or practice, you would also not be equipped to be a responsibility taker.

Second, it's possible for someone to be trained into automatic compliance with normative expectations without having the cognitive capacities to discern what value is served by particular practices, or to discern what, outside of the standard normative expectations, might serve to promote that value. Children, I assume, are to some extent accountability and compliance responsible before they are equipped to take responsibility.

Third, and more speculatively, someone could be a kind of visionary, keenly devoted to promoting some good, say, social justice, which some of our practices serve, but on whom social norms have little to no hold because the individual is invested in thinking for herself about how best to promote those goods. So, someone might be a responsibility taker but not compliance responsible within practices where social justice is at stake.

Finally, as I suggested earlier in this lecture, it's conceivable that there be normative communities in which *nothing* is normatively expected, but rather all is left to participants' election. Any dividing line between the normatively expected (what might be demanded by social participants) and the normatively optional (what is good but not demandable) is a social artifact. I see no reason to think that any possible society must operate with this distinction between the demandable and the not-demandable. A society of well-behaved agents might find that what needs to get done in fact gets done without normative expectations and their associated sanctions. Indeed, one might think that friendships and love-relationships largely operate this way, and possibly some religious communities.

That the three dimensions of being a responsible person can come apart in this way is, I think, a virtue of my account. The account captures default presumptions about fellow social practitioners in social life as we know it. Social engagement within the practices with which we are familiar presuppose all three dimensions of responsible persons. But the account also allows us to say why outsiders to our practices and young children should be engaged with as responsible persons rather than merely as objects to be managed. It does so by recognizing that some responsible persons may be merely accountability responsible because, not being part of our shared normative community or being very young, they can only be addressed as developable agents;

and older children, who may have reached the stage of being compliance responsible, may not yet have grasped the underlying goods served by norm-structured practices or have the cognitive capacities for discerning how those goods might be promoted in non-required ways. Of less importance, but still significant, I think we want any account of responsible persons to include societies composed of well-behaving individuals for whom having demandable normative expectations is unnecessary, and all that ought to be done can be left to voluntary election.

I'll close with a final virtue of the account of responsibility I've offered. I have not approached the topic of responsibility in the usual way by focusing on what is requisite for being responsible *for* particular actions. Instead, I've focused on having the valuable *status* "responsible person." That meant shifting from thinking of responsible persons primarily as beings who might deserve sanctions to thinking of them as beings *worthy* of responses befitting their valuable status. The shift is from thinking about what we as offended parties might be *entitled to do to* responsible persons to thinking about how we might *offend responsible persons* themselves by insultingly treating them as though they were not, or not fully, responsible persons.


References

- Alfano, Mark. 2021. "Towards a Genealogy of Forward-Looking Responsibility." *The Monist* 104 (4): 498–509.
- Baier, Annette. 1986. "Trust and Antitrust." *Ethics* 96 (2): 231–60.
- Beglin, David. 2019. "Two Strawsonian Strategies for Accounting for Morally Responsible Agency." *Philosophical Studies* 177 (8): 2341–64.
- Bero, Stephen. 2020. "Holding Responsible and Taking Responsibility." *Law and Philosophy* 39 (3): 263–96.
- Bicchieri, Cristina. 2005. *The Grammar of Society: The Nature and Dynamics of Social Norms*. New York: Cambridge University Press.
- . 2016. *Norms in the Wild: How to Diagnose, Measure, and Change Social Norms*. Oxford: Oxford University Press.
- Brennan, Johnny. 2021. "Recognition Trust." *Philosophical Studies* 178 (11): 3799–3818.
- Buchanan, Allen. 1996. "Perfecting Imperfect Duties: Collective Action to Create Moral Obligations." *Business Ethics Quarterly*: 27–42.
- Calhoun, Cheshire. 2021. "Appreciating Responsible Persons." In *Oxford Studies in Normative Ethics Volume 11*, edited by Mark Timmons, 9–28. Oxford: Oxford University Press.
- Carbonell, Vanessa. 2017. "Social Constraints on Moral Address." *Philosophy and Phenomenological Research* 98 (1): 167–89.
- Darwall, Stephen. 2006. *The Second-Person Standpoint: Morality, Respect, and Accountability*. Cambridge, MA: Harvard University Press.
- D’Cruz, Jason. 2019. "Humble Trust." *Philosophical Studies* 176 (4): 933–53.
- Eshleman, Andrew S. 2014. "Worthy of Praise." In *Oxford Studies in Agency and Responsibility, Volume 2*, edited by David Shoemaker and Neal Tognazzini, 216–42. Oxford: Oxford University Press.
- Fernández Castro, Victor and Elisabeth Pacherie. 2020. "Joint Actions, Commitments and the Need to Belong." *Synthese* 198 (8): 7597–626.
- Ferry, Michael. 2015. "Beyond Obligation: Reasons and Supererogation." *Royal Institute of Philosophy Supplements* 77: 49–65.
- Fingarette, Herbert. 1966. "Responsibility." *Mind* 75 (297): 58–74.
- Fricke, Miranda. 2007. *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford: Oxford University Press.
- Garfinkel, Harold. 1963. "A Conception of and Experiments with 'Trust' as a Condition of Concerted Stable Actions." In *Motivation and Social Interaction: Cognitive Determinants*, edited by O. J. Harvey, 187–238. New York: Ronald Press.
- . 1964. "Studies of the Routine Grounds of Everyday Activities." *Social Problems* 11 (3): 225–50.
- Gilbert, Margaret. 2008. "Social Convention Revisited." *Topoi* 27 (1–2): 5–16.
- Govier, Trudy. 1997. *Social Trust and Human Communities*. Montreal; Kingston; London; Buffalo: McGill-Queen’s University Press.
- Helm, Bennett W. 2012. "Accountability and Some Social Dimensions of Human Agency." *Philosophical Issues* 22 (1): 217–32.
- Herman, Barbara. 2007. *Moral Literacy*. Cambridge, MA: Harvard University Press.

- . 2013. “Morality and Everyday Life.” *Proceedings and Addresses of the American Philosophical Association* 74(2): 29–45.
- Holroyd, Jules. 2007. “A Communicative Conception of Moral Appraisal.” *Ethical Theory and Moral Practice* 10 (3): 267–78.
- . 2018. “Two Ways of Socializing Moral Responsibility.” In *Social Dimensions of Moral Responsibility*, edited by Katrina Hutchison, Catriona Mackenzie, and Marina Oshana, 137–62. New York: Oxford University Press.
- Jones, Karen. 1996. “Trust as an Affective Attitude.” *Ethics* 107 (1): 4–25.
- . 2017. “But I Was Counting on You!” In *The Philosophy of Trust*, 90–108. Oxford University Press, Oxford.
- Levy, D. K. 2015. “Assimilating Supererogation.” *Royal Institute of Philosophy Supplements* 77: 227–42.
- Macnamara, Coleen. 2013. “‘Screw You!’ & ‘Thank You.’” *Philosophical Studies* 165: 893–914.
- . 2015. “Reactive Attitudes as Communicative Entities.” *Philosophy and Phenomenological Research* 90 (3): 546–69.
- Martin, Adrienne M. 2013. *How We Hope: A Moral Psychology*. Princeton; Oxford: Princeton University Press.
- Mason, Michelle. 2017. “Reactive Attitudes and Second-Personal Address.” In *Ethical Sentimentalism*, edited by Remy Debes and Karsten R. Stueber, 153–70. Cambridge University Press.
- McGeer, Victoria. 2012a. “Civilizing Blame.” In *Blame: Its Nature and Norms*, edited by D. Justin Coates and Neal A. Tognazzini, 162–88. New York: Oxford University Press.
- . 2012b. “Co-Reactive Attitudes and the Making of Moral Community.” In *Emotions, Imagination, and Moral Reasoning*, edited by Robyn Langdon and Catriona Mackenzie, 299–326. New York; London: Psychology Press.
- . 2015. “Building a Better Theory of Responsibility.” *Philosophical Studies* 172: 2635–49.
- . 2019. “Scaffolding Agency: A Proleptic Account of the Reactive Attitudes.” *European Journal of Philosophy* 27 (2): 301–23.
- McKenna, Michael. 2012. *Conversation and Responsibility*. New York: Oxford University Press.
- Preston-Roedder, Ryan. 2013. “Faith in Humanity.” *Philosophy and Phenomenological Research* 87 (3): 664–87.
- . 2017. “Civic Trust.” *Philosophers’ Imprint* 17 (4).
- . 2018. “Three Varieties of Faith.” *Philosophical Topics* 46 (1): 173–99.
- Rawls, Anne Warfield and Gary David. 2005. “Accountably Other: Trust, Reciprocity and Exclusion in a Context of Situated Practice.” *Human Studies* 28 (4): 469–97.
- Shoemaker, David. 2007. “Moral Address, Moral Responsibility, and the Boundaries of the Moral Community.” *Ethics* 118 (1): 70–108.
- . 2019. “Hurt Feelings.” *The Journal of Philosophy* 116 (3): 125–48.
- . 2020. “Empathic Self-Control.” In *Surrounding Self-Control*, edited by Alfred Mele, 384–99. New York: Oxford University Press.
- Stohr, Karen. 2019. *Minding the Gap*. New York: Oxford University Press.
- Strawson, P. F. 2020. “Freedom and Resentment.” In *Freedom, Resentment, and the Metaphysics of Morals*, edited by P. F. Strawson, 107–34. Princeton; Oxford: Princeton University Press.
- Telech, Daniel. 2020. “Demanding More of Strawsonian Accountability Theory.” *European Journal of Philosophy* 28 (4): 926–41.


- . 2021. "Praise as Moral Address." In *Oxford Studies in Agency and Responsibility Volume 7*, edited by David Shoemaker 154–81. Oxford: Oxford University Press.
- Thomas, Laurence. 1989. "Trust and Survival: Securing a Vision of the Good Society." *Journal of Social Philosophy* 20 (1–2): 34–41.
- Trianosky, Gregory W. 1986. "Supererogation, Wrongdoing, and Vice: On the Autonomy of the Ethics of Virtue." *The Journal of Philosophy* 83 (1): 26–40.
- Vargas, Manuel. 2013. *Building Better Beings: A Theory of Moral Responsibility*. Oxford: Oxford University Press.
- Walker, Margaret Urban. 2006. *Moral Repair: Reconstructing Moral Relations After Wrongdoing*. Cambridge: Cambridge University Press.
- Wallace, James D. 2008. *Norms and Practices*. Ithaca, NY: Cornell University Press.
- Wallace, R. Jay. 1994. *Responsibility and the Moral Sentiments*. Cambridge, MA; London: Harvard University Press.
- Wolf, Susan. 1990. *Freedom within Reason*. New York: Oxford University Press.

R



COMMENTARIES

**Jules Holroyd
Heidi Maibom
Gunnar Björnsson**



4

NON-IDEALIZED SOCIAL PRACTICES: RESPONSE TO CALHOUN

Jules Holroyd
The University of Sheffield

4.1 Introduction

Across the three lectures, Cheshire Calhoun presents us with an expanded conception of morally responsible persons: the capacities that underpin moral responsibility, the responsibility-recognizing attitudes with which we respond to each other qua responsible persons, and the social practices in which the responsibility practices are embedded.

The conception of responsible persons that Calhoun develops is tripartite, with three dimensions of the responsible person brought to the fore: the accountability-responsible person, the compliance responsible person, and the responsibility-taking person. Each aspect of responsible persons has a role in our norm-structured social practices, and legitimizes different kinds of responsibility-recognizing attitudes. In particular, these attitudes are wider than those of blame and resentment that have occupied much of the philosophical landscape.

In broad brush strokes, the *accountability-responsible person*—person being a cross-temporal status (3–4)—has developable and sometimes realized capacities to understand, detect, and govern themselves according to normative expectations (11). When others have these capacities, it is legitimate to hold them accountable, and to respond with blame when they fail to meet our normative expectations. But our capacity for development also warrants our treatment of each other as educable, and as people from whom we, qua developable agents, can learn. The expression of the

reactive attitudes, when normative expectations have been violated, is partly how we learn and develop our capacities. But this conception also enables us to capture those people who are not yet part of our normative communities, but are nonetheless responsible persons, in light of their developable, sensitizable, capacities.

The *compliance-responsible person* not only has the capacities to meet normative expectations, but also has the normative literacy, in a particular context, to function within norm-governed social practices. The compliance responsible person, we presume, not only has the relevant normative literacy to be able to comply, but for the most part is disposed to do so (31). Thus the responsibility-recognizing (and crucially, not reactive) attitude of default trust characterizes our interpersonal interactions with compliance responsible persons (39–40). Moreover, with compliance-responsible persons, we can feasibly expect uptake and response in our communicative expressions of the reactive attitudes (41).

The *responsibility-taking person*, finally, has the capacity to grasp the point of normative expectations and values that structure our practices, and so to elect to do good things that are not normatively required of them (62). These capacities show up in our prospective responsibility-recognizing attitudes and practices—those of asking for volunteers, relying on non-obligatory assistance, asking for favors (62–63)—taking on responsibilities (67). This conception of the responsible person is embodied in the positive reactive attitudes of praise and gratitude (58).

Much theorizing about moral responsibility has been concerned with the accountability-responsible person; it is with this narrow focus, and the assumptions underpinning it, that Calhoun takes issue. Crucially, being a responsible person amounts to much more, on Calhoun's picture, than simply being liable to blame—as the more expansive tripartite conception, and its associated responsibility-recognizing attitudes and practices, indicate. Calhoun's aim is to develop a “much richer, more complexly three-dimensional conception of responsible persons” (8) than is ordinarily assumed; the idea is that this picture can make better sense of the value of the status of “responsible person,” and the value of the social practices in which such persons are embedded (3–6).

In Section 1.2, I will focus on the assumptions that Calhoun rejects, and endorse these moves. In Section 1.3, however, I will raise some concerns about the somewhat idealized view of social practices that takes the foreground in Calhoun's lectures. I will suggest that when we attend to some less idealized aspects of our social practices, we find a more equivocal value in the responsibility practices. In Section 1.4, I also argue

for a messier picture of the tripartite aspects of responsibility, and their responsibility-recognizing attitudes in particular, and show how this interacts with Calhoun's critique of how the positive reactive attitudes have been analysed. I argue that paying attention to less idealized practices renders visible the role of the positive reactive attitudes in exerting social pressure—sometimes in valuable, constructive, ways, but also in relation to oppressive norms and values. I see Calhoun's more idealized project as articulating the potential value of being a responsible person, and of our practices; my project, taking a non-idealized starting point, is to articulate the conditions in which those values are realized or undermined. This can in turn inform the articulation of the contours of the practices and concepts themselves.

4.2 Rejecting Four Assumptions

As is already clear from the brief characterization above, Calhoun's account of morally responsible persons expands the scope of our theorizing about moral responsibility: it presents a wider range of capacities relevant to being a responsible person (in the three different dimensions of responsibility), and it therefore accommodates a wider range of responsibility-recognizing attitudes. In so doing, Calhoun moves beyond a pervasively-found approach to moral responsibility that is shaped around the following four assumptions (paraphrasing from pages 7–8):

Assumption 1: The only capacities relevant to being a responsible person are those that warrant blaming attitudes and holding to account for wrongdoing (absent excuse or exemption).

Assumption 2: What it is to see someone as a responsible person is to see them as someone whom we hold accountable to normative expectations.

Assumption 3: All responsibility-recognizing attitudes are reactive: to blameworthy failures to meet those normative expectations, or credit-worthy exceedings of normative expectations.

Assumption 4: Our responsibility practices are exhausted by our practices of holding others accountable to normative expectations (for failing to meet, or exceeding, them).

Calhoun develops the tripartite conception of moral responsibility to show where each of these assumptions is mistaken. In the following sections I reconstruct her claims in brief.

4.2.1 *Rejecting Assumption 1*

Assumption 1: The only capacities relevant to being a responsible person are those that warrant blaming attitudes and holding to account for wrongdoing (absent excuse or exemption).³⁶

Theorists have focused on the “capacity to understand normative concepts, to detect normatively relevant considerations, and to deliberate and govern one’s actions in light of normatively relevant considerations” (12) as relevant to responsibility. On this assumption, when someone is not liable to blame for their wrongdoing it is because they lack the relevant capacities for moral responsibility.

Calhoun argues that being accountable involves having the “developable capacity to become sensitized to a wider range of normative reasons” (17).³⁷ This is a cross-temporal capacity: we don’t lose it when we happen not to be sensitive to certain normative considerations, through error or ignorance—nor even through malice. We retain this capacity, and our status as responsible persons, even if we are not, on occasion, liable to blame (say, because of our unfamiliarity with a local normative context).³⁸ It is a mistake—an insulting one—to suppose that agents who are not liable to blame are not morally responsible persons, lacking even the capacities to develop our normative sensibilities (comparable, in respect of moral responsibility, to cats).

Moreover, being a person with compliance responsibility, and who takes responsibility, requires distinctive capacities, beyond those involved in the basic “ground floor” accountability responsibility. It involves, respectively, the capacity for normative literacy in a particular context, and being disposed to comply; and having the capacity to grasp the point of certain normative expectations, such that one can sometimes elect to promote the values that those norms underpin. These capacities clearly go beyond the narrow set of capacities which theorists have focused on as relevant to accountability responsibility.

^{36.} One might think this clashes with Assumption 3, but the assumption is that the same capacities are engaged in both blameworthy failures and praiseworthy exceedings of expectations.

^{37.} Calhoun is here drawing on McGeer (2019).

^{38.} The foil here in Calhoun’s discussion is the position defended by Vargas (2013)—circumstantialism—whereby a person’s status as responsible shifts depending on their realizing the relevant capacities across circumstances. For a discussion and critique of Vargas’s circumstantialist claims, see also Holroyd 2018.

4.2.2 *Rejecting Assumption 2*

Assumption 2: What it is to see someone as a responsible person is to see them as someone whom we hold accountable to normative expectations.

Beyond seeing other morally responsible people as those we hold to normative expectations, then, we also see responsible persons as those who can hold us to account, as well as persons who are educable, and from whom we can learn. But we also see others as responsible when we take on the default position of assuming that they will comply with normative expectations, and that they have the normative literacy and dispositions to do so (i.e., compliance responsibility). When we treat others as persons capable of electing to do good, and organize our practices around this—by assuming that there will be volunteers, soliciting donations, expecting, at least sometimes, favors—we also treat people as responsible persons, who are capable of taking responsibility.

4.2.3 *Rejecting Assumption 3*

Assumption 3: All responsibility-recognizing attitudes are reactive: to blameworthy failures to meet those normative expectations, or credit-worthy exceedings of normative expectations.

Whilst theorists have focused on reactive attitudes such as blame and resentment, in response to wrongdoing, these depict only a partial view of what Calhoun calls the “responsibility-recognizing attitudes.” These attitudes extend beyond backward-looking resentment to include our default (and therefore less vividly present to conscious awareness) attitudes of trust in others’ normative compliance (11, 38–40), and a tendency to give others the benefit of the doubt. If one were “immediately accused of norm violation, rather than given the benefit of the doubt,” one would not be afforded the presumption of being a compliance-responsible person, Calhoun claims (45). Our responsibility-recognizing attitudes also extend to include our prospective attitudes of hope (that others will elect to do good, or promote some value) (63) as well as the reactive attitudes of appreciation, praise, and gratitude when they do so (exercising their capacities qua responsibility takers) (62). The more expansive tripartite account of the responsible person and their capacities thus delivers a more expansive set of corollary responsibility-recognizing attitudes.

4.2.4 *Rejecting Assumption 4*

Assumption 4: our responsibility practices are exhausted by our practices of holding others accountable to normative expectations (for failing to meet, or exceeding, them).

And likewise, we thus have a richer view of the social practices in which our status as responsible persons matter, and in which the responsibility-recognizing attitudes are present or activated. Whereas theorists have focused narrowly on “responsibility practices” of holding each other to account, on Calhoun’s view we now see that it is legitimate to consider the everyday practices in which we rely on each other’s default trust as part of the social practices afforded by, and available to, responsible persons. And, the wide range of practices and activities that depend on persons taking responsibility—electing to do good, beyond what is normatively required of them, including in workplaces, community groups, families, friendships—are also brought within the remit of responsibility practices. Thus, once again, the value of the status of “responsible person”—given its significance in so many of our activities—is affirmed.

4.3 Methodology: The Social Practices

Calhoun’s lectures follow the broadly Strawsonian method of looking at our practices, which she characterizes as follows:

If we want to understand responsibility, we need to begin by reflecting on what social life with others is like, rather than beginning from metaphysical views about contra-causal freedom or determinism. (51)

Rather than retreat to such “panicky metaphysics,” to use Strawson’s term (1962, 25), we should start instead with “the facts as we know them,” and build our conception of the responsible agent or person from these facts. But, Calhoun asks, which facts? One key methodological development Calhoun recommends is to start not with a narrow conception of our responsibility practices—our intuitive understandings of what happens when we hold others accountable for failing to live up to normative expectations. Rather, we should attend to our *social practices* more broadly (8). Thus her rejection of Assumption 4, as we have seen. I am generally favorably disposed to this approach. But, we might ask, which social practices?

4.3.1 *Social Hierarchies and Default Status*

What motivates this question is a worry about the rather idealized conception of social practices that Calhoun foregrounds in her lectures. These are social practices

in which, Calhoun argues, the default attitude is trust in others—both to be capable of meeting normative expectations, and of having the normative literacy and dispositions that enable compliance with these normative expectations.

Calhoun does mention dynamics of social power, and acknowledges (in her discussion of Carbonell’s work) that it “is an unfortunate feature of hierarchical societies that the moral power to hold to account gets unequally distributed” (18).³⁹ And, as she writes of the kinds of insult that one might experience on the basis of one’s social group membership: “The insult [of not being trusted to comply with normative expectations] that Black individuals receive when they are treated with suspicion in all sorts of social practices—driving while Black, shopping while Black, and so on” (44).

But since social power and hierarchy are not incidental, but systematic, and since social identities such as race, gender, and class are not momentary but tend to track us across many (if not all) social contexts, these departures from the default will not be incidental aberrances, but rather systematic failures.⁴⁰

As Ciurria writes, in challenging the assumption about our default attitudes to others (an assumption I myself have previously made):

In conditions of structural injustice, many people do not enjoy the privilege of being presumed responsible until proven otherwise. Disabled folks, for example, are generally presumed non-responsible until they prove themselves capable of “adult moral reasoning,” and this bias permeates the philosophical literature (including Strawson’s work). Black people are, as critical race theorists have shown, presumed subhuman and uncivilized as a group, which explains why they are overincarcerated, overpoliced, and disproportionately subjected to family separation. [This] assumption [of responsible until proven otherwise] doesn’t square with the reality that being presumed responsible is a privilege, generally withheld from disenfranchised groups. (2021, 176)⁴¹

³⁹. And, these dynamics have been prominent in her earlier work, e.g. “Responsibility and Reproach” (1989).

⁴⁰. See, however, Glasgow (2007) on race not traveling.

⁴¹. The claim that disabled people are “generally presumed non-responsible” needs qualification: this will be true for some ways of being disabled and not others.

Starting from the social facts about social power and social hierarchy as we know them, then, challenges the idea that our default attitudes are respectful ones that recognize others' status as responsible persons. Calhoun claims that "the status 'accountability responsible person' travels" (21); that is, whether or not one is an accountability responsible person does not depend on the particularities of one's normative context and community. Ciurria's remarks challenge this claim. Whilst the capacities that underpin this status might travel, being recognized as having such a status is not context-independent; it will depend on the social hierarchies operative in any particular social practice. For example, a Black person might be recognized as having the status "responsible person" within their family or their community organizing group; but not be accorded this status in other contexts, such as in interactions with law enforcement, or child welfare (or "family regulation," in Roberts' terms) systems (see e.g. Roberts 1997; 2022).

Likewise, facts as we know them about social hierarchy and oppression challenge the idea that the attitude of default trust (characteristic, on Calhoun's view, of regarding others as compliance responsible) is extended to all; nor is it held by all. Consider the scenario described by Elizabeth Anderson:

One late night in 2007 I was driving in Detroit when my oil light came on. I pulled into the nearest gas station to investigate the problem when a young black man approached me to offer help. "Don't worry, I'm not here to rob you," he said, holding up his hands, palms flat at face level, gesturing his innocence. "Do you need some help with your car?"... This encounter illustrates the public standing of racial stereotypes as default images that influence the interactions of black and white strangers in unstructured settings, even when both parties are prepared to disavow them. A little ritual must be performed to confirm that both parties do disavow the stigma, so that the cooperative interaction may proceed. (2010, 53)⁴²

Note that the lack of trust assumed in this scenario does not feature as an incidental insult; rather it is the default setting that both individuals have to strive to overcome. Racist stereotypes condition the very (non-)existence of the default trust. The would-be helpful stranger could neither assume that Anderson trusted him, nor trust that she wasn't making racist assumptions and wouldn't respond with aggression to his approach.

⁴². See also Haslanger 2015 for a discussion of this case in terms of schemas and social meaning.

The attitude of default trust may also be eroded by particular experiences of oppression. After a man raped her, Ellie Campbell described being confronted by an erosion of trust:

I'd felt safe around men for years ... and then suddenly I was like how am I in this place where I'm really scared and really volatile, like something bad could happen at any moment.⁴³

Elsewhere, she writes “The rape ... was trying to make me fearful of sex, and of men” (2022). Default trust is eroded by misogyny—by sexual violence—and, it is implied in Campbell’s writing, by the victim blaming attitudes she experienced afterwards: “How can someone I know, and should trust, suggest that my obsession with men got me raped?”

The methodology of starting not just with social practices, but non-ideal social practices, characterized by social hierarchy and attendant stereotypes and oppressions, reveals a less rosy picture of those social practices, and the default attitudes that characterize them. Rather than default attitudes of respect for status as responsible, and trust in compliance with norms, we find instead both disrespectful and dehumanizing stereotypes, and reasonable withdrawals of trust in the face of oppression. This prompts us to consider a more nuanced question about the relationship between our responsibility practices and default attitudes. Rather than assuming that default trust characterizes our social practices, we should ask: to what extent can our social practices and responsibility practices withstand the erosion of default trust?

4.3.2 *Oppressive Norms and the Value of the Status “Responsible Person”*

A second idealization that Calhoun makes concerns the value of the status “responsible person.” We can see that the status of being a responsible person is valuable, Calhoun observes, by noting that *not* being accorded that status is an insult. Developing an account that recognizes this value is important. Accounts which focus narrowly on accountability for failure to meet normative expectations are not well placed to accommodate this value. As Calhoun puts it, “Who, after all, wants to have demands pressed upon them and be blamed and resented?” Why would liability to blame be something that we would want and value?

⁴³ Timestamp 15:30 of the Today in Focus podcast “Why Comedian Grace Campbell Refuses to Be Silenced about Sex” (2022).

Instead, a conception of responsible persons that recognizes that it involves *developable* capacities, the ability to educate and learn from each other, the basic trust that comes with being treated as compliance responsible, and the positive prospective and reactive attitudes that come with taking responsibility, better comports with the value of such a status.

I do not here want to question Calhoun's claims that such a status has value. However, I want to point to the various burdensome aspects of participation in the practice—beyond potential liability to blame—that are brought into view when we consider non-ideal social practices. These burdens may temper the value of the status. As such, if the value of the status “responsible person” is to be adequately realized, these aspects of the non-ideal social practice that threaten to undermine its value must be recognized and addressed.

As we have already seen, Calhoun acknowledges that the ability to “cash in” (as it were) on the value of the status of responsible person might differ depending on one's social position and social context. Calhoun points to Carbonell's work on the compromised ability of members of a family of Hmong refugees, in the U.S. in the 1980s, to hold others to account—unable to challenge moral decisions of others either in the moment or in calling for recognition of and repair for wrongs. In such cases, “the marginalized person's moral demand is ignored, misinterpreted, underestimated, rejected, or silenced” (Carbonell 2019, 178; cited by Calhoun on page 19 of this volume).

But even when one's status as a responsible person is recognized, this may be burdensome in various ways. Responsibility-recognizing attitudes may be directed towards responsible persons in accordance with oppressive norms, for example. This is true of both blame and praise (and plausibly other responsibility-recognizing attitudes also). Consider Ciurria's remarks on how women are held accountable for being victims of sexual violence: in strands of contemporary discourse it is “women's ‘risky decisions’ that expose them to injury and violence ... [this] discourse illicitly blames women for men's transgressions (e.g., rape [...]), thereby enforcing patriarchal relations” (2020, 5–6). When the status of “accountability-responsible person” is recognized, it can license distinctive ways of perpetuating oppression via the responsibility-recognizing attitudes, by erroneously allocating blame to women, and enabling men who perpetrate violence to escape responsibility.

Praise, too, can perpetuate oppressive norms. Aubrey Gordon (2020), formerly writing under the pseudonym “Your Fat Friend,” remarks on her objection to being

praised for her “bravery” in wearing what she ultimately describes as a rather “unremarkable, standard” dress:

I was only *brave* if my body was meant to be a source of shame, something to be shut away, covered up, rarely seen and never discussed. And she [the appraiser] simply couldn’t conceive of someone with a body like mine daring to get dressed, daring to be seen, daring to show up in the same places as someone with a body like hers.

There are different reads of this case: on one, the appraiser is really just using the appraisal as a cover for insult; on a second, there is a genuine intention to praise (and no intention to insult), shaped by internalized oppressive norms; a third might involve someone close enough to the *praisee* to understand well the challenges and obstacles they might face, and who is able to praise in a way that does not express or imply oppressive norms (for example, knowing that the person has endeavored to ignore oppressive norms of appearance). It is important to my point that the second interpretation is the right one in this particular case; this understanding is supported by Gordon’s critical interpretation of the interaction (in other scenarios, other interpretations may better fit the particulars of the situation). Thus understood, such praise itself serves (albeit unintentionally) as an insult, as expressing harmful norms, as communicating the expectation of body shame, alongside the praise for bravery. Thus, being recognized as, and interacted with as, a responsible person in our social practices—even in the more capacious set of practices to which Calhoun draws attention—may sometimes have equivocal value for those whose participation in those practices is shaped and constrained by social hierarchy. The point isn’t that there is no value in the status, but that the burdens attendant on the status in non-ideal social practices are revealed if we look more closely at the non-ideal versions of those practices. Being recognized as a responsible person, and being responded to with responsibility-recognizing attitudes, is perfectly compatible with, and can itself consist in, various kinds of responsibility-related insult. Again, whilst we can accept that there is value in having the status of responsible person, we ought to ask which social practices best realize that value, and which undermine it.

4.3.3 *Tripartite Responsibility and the Responsibility-Recognizing Attitudes*

A further idealization we find in Calhoun’s work concerns the relationship between the responsibility-recognizing attitudes and their fit in relation to the three aspects of responsible persons. In her presentation of the claims, the picture we get (the one I endeavored to accurately reconstruct in Section 1.2) is one in which there is a neat alignment between the responsibility-recognizing attitudes and the aspects

of responsibility. Accountability responsibility sometimes renders apt attitudes of blame, but also hopes of education; compliance responsibility renders apt the attitudes of basic trust, predictive expectations for compliance which lead us to give people the benefit of the doubt; and persons who take responsibility are aptly the recipients of positive reactive attitudes: prospectively, hope; reactively: praise, gratitude.

I don't here challenge the tripartite model of responsible persons. Rather I want to point to instances in which, in fact, in our non-ideal social practices, there is a degree of messiness worth attending to. The point of these observations is not solely to depict a more complex view of those practices, whereby the attitudes and their apt targets are misaligned, if we take a different methodological starting point. My aim is also to bring out some insights that can inform our models of the reactive attitudes, in Section 1.4.

4.3.3.1 Praising Erroneously as Elective

In Calhoun's model, the positive reactive attitudes such as praise and gratitude are apt responses to responsibility-takers, who elect to do good. There are instances in which the expression of positive reactive attitudes appears problematic precisely because it treats certain good actions as elective, rather than normatively expected or required. Consider the case of the "Daddy dividend," introduced by Khader and Lindauer:

We coined the term "daddy dividend" on one of the many days a stranger on the subway told Matt (a white man) that he was the "best daddy ever." The thing he had done to receive this accolade was to wear his baby in a carrier, and perhaps not seem utterly miserable doing so. By contrast, Serene (a brown woman) has never been told by a stranger that she's the best mom ever, or even a decent one. (2020, 6)⁴⁴

Elsewhere I have teased out some of the problematic dynamics here: the distribution of praise, and the stereotype-informed low expectations of men expressed in such asymmetrical praise (Holroyd 2021). One further way of capturing what is going on here is that, in praising such parenting by fathers, it treats it as elective, rather than

⁴⁴ In fact, this case and my discussion of it has considerable resonance with Calhoun's observations, more than four decades ago, on how oppressive norms distort the reactive attitudes: "Sensitive to the social determinants of oppression, women often feel grateful when husbands volunteer to babysit ... [where oppressive contexts distort the logic of moral language:] No wonder, then, that women have trouble sustaining their sense of what is owed them and find themselves feeling grateful when given their due" (1989, 403–404).

what should be normatively expected of fathers.⁴⁵ In Calhoun's terms: if the positive reactive attitudes are directed towards the *optional promotion* of values (namely, the pursuit of values beyond what is normatively required of us), then such use of praise depicts as merely optional the involvement of fathers in basic parenting.

4.3.3.2 Default Trust for Violators of Normative Expectations

In Calhoun's picture we see that those treated as compliance-responsible—as beings who for the most part understand and are disposed to comply with normative expectations—are regarded with default trust. Moreover, this trust dictates that we should in general give responsible persons the benefit of the doubt:

Responsibility practices include ... giving the benefit of the doubt when those presumed to be compliance responsible appear to be violating normative expectations. There must be some misunderstanding that could be resolved, or it's just a joke, or there must be some excuse. (45)

But this responsibility-recognizing attitude can be distorted by oppressive norms. In the context of misogyny, for example, this attitude—apt for those we assume to be disposed to comply—is afforded even to those who have demonstrated themselves to be in violation of normative expectations. As Manne writes, the disproportionate sympathy extended to male perpetrators, and the unwillingness to dislodge the assumption that they are really “good guys,” means that even when there is clear evidence of normative expectations being violated—of a serious crime being committed—men are given the benefit of the doubt. This is what we saw in the case of Brock Turner's charge and ultimate conviction for rape, as Manne writes:

Despite the fact that Turner was caught in the act of violating Miller behind a dumpster (by two Swedish graduate students, who performed a citizen's arrest), many people expressed skepticism that Turner could possibly be a rapist. ... One of his friends opined that Turner's crime was “completely different ... That is a rapist. I know for a fact that

⁴⁵ In discussion, a number of people have also observed that there can be something patronizing about praising fathers for doing their basic roles in parenting: the praise seems to treat fathers as if their sensibilities qua parents are in need of encouragement, nurturing and honing—that their skills and moral characters as fathers are still in development. It infantilizes and patronizes fathers when they are praised for doing basic parenting. I think this is also a feasible interpretation, and one that trades on the idea that praise is sometimes aptly directed towards those who have merely developable capacities—a point I argue for in Section 1.4, below.

Brock is not one of these people,” she [a friend] wrote in a statement attesting to his good character. (2020, 20)⁴⁶

Giving the benefit of the doubt, Manne writes, can “imaginatively turn other crimes, such as rape, into mere misunderstandings and alcohol-fueled mishaps” (20). Defaulting to giving the accused the benefit of the doubt makes it harder for victims to report, and be heard, in pursuing justice.

To be clear: as I understand her, Calhoun isn’t claiming that we only give the benefit of the doubt to compliance-responsible persons, nor that it is never fitting elsewhere, nor that it might not be misplaced (as it certainly is in Turner’s case). We might give—and perhaps might be required to give—the benefit of the doubt to family members, and close friends, for example. And we might give it when we are required not to. My point is that in social practices characterized by oppression, we see the deployment of responsibility-recognizing attitudes apt for compliance-responsible persons towards those who have violated normative expectations. This misogynistic deployment of responsibility-recognizing attitudes serves to vindicate oppressors, and thereby enforce oppression and silence victims.

4.3.3.3 Weaponizing Taking Responsibility

On Calhoun’s view, when we treat others as beings who will *take responsibility*—who will go beyond what is normatively expected of them—we respect them as persons who are “capable of and sometimes disposed to take responsibility by electing to promote good ends in ways that aren’t normatively expected” (62). This aspect of our social practices “[makes] use of a set of prospective responsibility-recognizing practices: asking for volunteers, making appeals, issuing directed invitations and requests, asking for favors or help, relying on those who have responsibilities to look after some good to do more than the bare minimum” (62). Such interactions legitimate the prospective attitude of hope: “It is a hope that in fulfilling their assigned responsibilities, they will value the goods served by those responsibilities enough to do more than is minimally required” (63).

And yet consider the case, reported on by the BBC (Titheradge 2022), of an eight-year-old child at an after-school club in the UK, who was expected by members of staff to clean up after his five-year-old sister had soiled herself. The BBC reports: “Her mother said the incident was racist and an example of ‘adultification’, where young black children are perceived as older than they are. ... ‘How on earth do they

⁴⁶ Manne describes this case alongside a number of others, showing the pattern, rather than anomalous nature, of such attitudes.

think that is acceptable?’ the mother said. ‘They would never possibly ask a white child to do that’.⁴⁷

One reading of this case is that the child was wholly inaptly—and as a result of racist stereotypes—expected to *take responsibility* (i.e., to do something that should not be normatively expected of him). What was being asked of the child traded on the assumption that he cared for his sibling such that he would ensure her well-being, promote her interests. And yet there is nothing hopeful about this: he is being asked to take responsibility both on the basis of racist assumptions, and in a way that enables those who should in fact have responsibility to avoid doing the bare minimum.⁴⁸

Again, the point is not that being a person who can take responsibility—and being treated as such—is not ever valuable. It is that when we start our theorizing from a different understanding of “the facts as we know them,” we see aspects of our responsibility practices—including the expanded set of practices that Calhoun articulates—that are weaponized for oppressive ends. Accordingly, we should ask about the contexts in which their value can be properly realized.

4.3.4 *Summary*

Let us now recap, and identify the key points of agreement, disagreement, and development of Calhoun’s claims. First, I argued, contra Calhoun, that it is a mistake to see the status of accountability-responsible as traveling, and that the trust that is apt for compliance-responsible persons is not reliably a default. Whether the status, and the attitude, is present or absent depends (inter alia) on the norms and stereotypes operative in a social context.

Next, I agreed with Calhoun that there is value in having the status of responsible person, but I drew attention to the distinctive burdens beyond liability to blame that can be attendant, for some, upon being a participant within the practice: in particular the use of the responsibility-recognizing attitudes (both positive and negative) to enforce oppression.

Then, I suggested that in our non-ideal practices the responsibility-recognizing attitudes may not find their apt targets. Positive reactive attitudes might be directed

^{47.} Whilst the interaction is described as one in which the child was “asked” to perform this care, the report aptly, because of the power dynamic and circumstance, describes him as being forced to do so. Thanks to Jim Chamberlain for drawing my attention to this example.

^{48.} Consider also the phenomenon of “mental labor” whereby women in heterosexual families are assumed to take responsibility for certain organizational aspects of family life.

towards what should really be normatively expected, implying instead its status as normatively optional. Attitudes of trust, and in particular of giving others the “benefit of the doubt,” can be distorted by and enforce oppressive norms. And treating another as a responsibility-taker can be underpinned by oppressive stereotypes, enforce oppression, and enable others to renege on their responsibilities.

Why draw attention to these depressing features of our responsibility practices? My discussion here is not intended solely to point out that there is a different, and less optimistic, perspective on the social practices—“the facts as we know them”—from which we might start our thinking about moral responsibility. Rather, as we will see in Section 1.4, I will argue that, when we start our thinking from non-idealized social practices, in which our responsibility-recognizing attitudes are shaped by and attuned to oppressive social dynamics, we render visible important and otherwise easily overlooked features of those practices. I make this claim in relation to the positive reactive attitude of praise, in particular. In doing so I critique Calhoun’s concerns with existing analyses of praise, and articulate a facet of the reactive attitudes most visible in our non-ideal practices. However, this feature is an important component of what the reactive attitudes do, and so should feature in any complete account of these attitudes. Recognizing this feature presents a more ambivalent picture of the value of our responsibility-recognizing attitudes and practices, and explains why we should attend carefully to the justification and expression of the positive reactive attitudes, and the contexts in which their value can be more adequately realized.

4.4 Praise and Pressure

In this section, I consider Calhoun’s remarks about the positive reactive attitudes, and praise in particular, and suggest we can respond to her critiques of existing analyses of praise by examining more closely the non-idealized social practices.

4.4.1 *Accountability Responsible Persons as Apt Targets of Praise*

Recall that on Calhoun’s tripartite conception of responsible persons, the positive reactive attitudes such as praise and gratitude are apt responses to responsibility takers, who elect to do good. As Calhoun puts it: “Just as basic trust recognizes others as compliance responsible persons, gratitude, appreciation, and the like recognize actions that befit responsibility taking kinds of beings” (58). As we have seen, though, basic trust is neither always, nor only, afforded to compliance-responsible persons. And, likewise, I think it is a mistake to suppose that there will be—even ideally—alignment between responsibility-takers and the positive reactive attitudes. In particular, I want to suggest that we should reject Calhoun’s claim that those who are (merely) accountability-responsible, and not yet responsibility-

takers, are not also apt candidates for praise. In these lectures, Calhoun implies this commitment in her framing: that an articulation of responsibility-taking provides a way of “accommodating the positive reactive attitudes” (52). The commitment is made more explicitly in her 2022 paper:

Negative reactive attitudes, like resentment, indignation, and guilt recognize their targets as the kind of being who can be held accountable for normative failures and is responsive to actual failure. ... Positive reactive attitudes, like appreciation and gratitude, recognize their targets as the kind of being who can take responsibility for doing what could have been blamelessly avoided and respond to the actual taking of responsibility. (2022, 24)

The point of positive reactive attitudes, such as appreciation and gratitude, “is to get people to see themselves as someone who is seen by others as having taken responsibility and to elicit the feeling of being appreciated” (2022, 26).

My claim is not that the positive reactive attitudes do not ever have these targets, but that they need not: the alignment between responsibility-taking persons and positive reactive attitudes is, even in principle, a mistake. To see this, return to the idea of accountability-responsible persons. Consider that the important expansion of how we conceive of accountability, for Calhoun, is that it consists in “a developable capacity to expand [our] sensitivity to normative reasons” (21). At this point, Calhoun is drawing on McGeer’s (2019) work. But unlike McGeer, who emphasizes the importance of persons being sensitizable to moral reasons in order to justify *blaming* them, Calhoun wants to develop this idea in order to explain how persons can be accountability-responsible whilst *not* being liable to blame. This difference is not emphasized in Calhoun, but it is important: on McGeer’s view, (warranted) blame is one of the tools available to us through which we sensitize others’ moral sensibilities (2019, 314). Yet Calhoun is aiming to explain the importance of treating each other as accountability-responsible—as sensitizable and with capacities to be developed—even when we are not yet liable to blame. Therefore, it seems all the more important (and heartily resonates with Calhoun’s rejection of Assumption 3) to acknowledge the role of other interactions and responsibility-recognizing attitudes involved in developing our capacities.

One alternative way in which the relevant capacities may be developed is via praise, which could be a form of feedback that affirms and sustains responsible

persons' developing moral sensibilities. Various strands of empirical research have focused on the ways in which praise can communicate important information about competence, increase motivation, and foster the development of one's capacities (see summaries in Soenens and Vansteenkiste 2020, Corpus and Good 2020, and Eskreis-Winkler and Fishbach 2020).⁴⁹ For example, Eskreis-Winkler and Fishbach remark on the finding that non-experts engage more—are more attentive, demonstrate more learning—when they get positive rather than negative feedback (2020, 51). Moreover, such praise need not be merely instrumental. Coates (2019) emphasizes the ways in which even those whose moral capacities are very much in development (e.g. small children) can demonstrate genuine achievements that render praise fitting.

This evidence supports Calhoun's point that the excessive focus on blame, resentment, and other negative reactive attitudes is mistaken. But, if praise can be (at least sometimes) aptly used as a capacity developing responsibility-recognizing attitude, then this should inform our evaluation of Calhoun's critiques of some communicative models of praise. For example, Calhoun objects to recent communicative accounts of praise that have focused on the *recognitive* aspects of praise (Macnamara 2013), or that frame it as seeking uptake from the praised person in a form of *shared valuing* of the good thing done (Telech 2021).

Such views model praise as a kind of communicative act, directed towards the praiseworthy person, that seeks uptake in a certain kind of attitude (acceptance of credit), emotion (directed pride), and behavior (expressions such as "You're welcome"). But, Calhoun suggests, since the person praised is often aware that they have done a good thing, and is often praiseworthy precisely because of their having done the good thing for the right kinds of reasons, it is odd to think of the point of communicative praise as pointing to the relevant properties (in contrast, it makes sense to draw the attention of wrong-doers to the relevant features for which they are blameworthy). As Calhoun puts it: "Expressions of gratitude and praise do something. But what they don't generally do is focus the target's attention on something they weren't already aware of in the first place" (58). And elsewhere, in relation to the positive reactive attitudes of appreciation or gratitude: "Any message whose gist is 'You did a good thing' simply affirms what the appreciative or grateful person must assume the target already knows" (2022, 25). Likewise, the idea that emotional uptake is sought by praise—in the praised person feeling pride—is also misplaced, Calhoun

⁴⁹ Though some studies have also highlighted ways in which some kinds of praise—depending on what it expresses—can be problematic: decreasing motivation, self-esteem, and perceptions of competence. See discussion in Corpus and Good 2020.

claims. Rather, the grounds for pride already exist (in the praiseworthy thing), and so precede the expression of praise: “The agent’s emotional ‘uptake’ typically precedes communicative exchanges; it is not reasonably an aimed-for effect of them” (58). No behavioral response, either (“You’re welcome”), seems required for the successful communication of praise, Calhoun claims (see pages 57–58).⁵⁰

If praise can be an effective and apt expressive tool for developing responsibility-relevant capacities, however, some of these observations require qualification. Where it is used to “scaffold” (in McGeer’s (2019) terms), we might think it particularly important that the praise draws attention to the good thing done, that it encourages pride, and that it therefore involves the appropriate kinds of cognitive and emotional uptake. Praising a person with merely developable capacities need not affirm what the praiser must assume the target already knows; the praise may be an important way of drawing attention to what has been done well as a means of developing the capacities of the praised person. Thus, if Calhoun is to expand the range of accountability-responsible persons to include those with merely developable capacities, it also makes sense to see praise as an apt response not only for those agents who can take responsibility, but also those whose capacities are not yet fully realized. The models of praise that seem best adapted to accountability responsibility (whereby praise seeks recognition of the good or shared valuing, encourages pride) may well be fitting here.⁵¹

Still, we might agree with Calhoun that these models are at best a partial view of praise, and that where praise is directed towards responsibility-takers—those with well-developed capacities who have elected to promote some value—these models of praise are less fitting. I do not disagree that a different model of responsibility may be a better fit in other instances of praise. However, in line with Calhoun’s exhortation to consider the *social practices* in which we interact as responsible persons, I want to fill out a model of what praise does by drawing on some of the instances of praise in non-ideal social contexts that we have seen. These cases provide evidence of the role that praise can play in exerting social pressure and enforcing social norms. Any complete account of praise should accommodate this role. This role, again, can rationalize the need for recognitive and emotional uptake of praise.

4.4.2 *Praise and Pressure*

⁵⁰. See also the expression of these lines of critique in Calhoun 2022, 24–26.

⁵¹. One need not think that the praise in these contexts is purely instrumental; one can view the target as deserving, and praiseworthy, given appropriately contextualized (to their stage of development) expectations.

As we have seen, Calhoun worries that (contra the claims of Macnamara 2013 and Telech 2021) there is little point in seeking recognitional or emotional uptake from responsibility-takers. However, I want to suggest that the purpose of praise in doing these things can be seen if we consider firstly one of the functions of the reactive attitudes—and specifically the social role of praise in non-ideal contexts—and secondly the role of praise as an expression that targets audiences other than its direct second-personal addressee.

4.4.2.1 Affirmation and Presupposition

Our reactive attitudes (the responsibility-recognizing attitudes that react to specific deeds) have a role in our social practices beyond simply responding to morally responsible persons' doings. This role has been articulated by Wallace, who writes that when we express the reactive attitudes:

We are demonstrating our commitment to certain moral standards, as regulative of social life. ... [The reactive attitudes] can be seen to have a positive, perhaps irreplaceable contribution to make to the constitution and maintenance of moral communities: by giving voice to the reactive emotions, these responses help to articulate, and thereby to affirm and deepen, our commitment to a common set of moral obligations. (1994, 69)⁵²

Wallace is here talking about blame specifically, the role of blaming responses in relation to violations of moral obligations, and commitments to moral standards.⁵³ But other reactive attitudes, including positive reactive attitudes, can play this role in expressing and affirming the values of a community. How might such expressions play this role—how do they affirm and deepen commitments to those values? By looking at instances of praise in non-ideal contexts, we can see one mechanism by which praise might serve this function, and how, in doing so, it is important that it seeks the kind of uptake that Calhoun doubted the purpose of.

First, recall the instances of problematic praise I have introduced: Gordon is praised for her supposed “bravery” in overcoming imputed body shame (Section 1.3.2); Lindauer is praised as “the best daddy ever” for performing basic parenting tasks

⁵². See also Macnamara 2013, 894; and Helm 2019 on the social function of the positive reactive attitudes (with a focus on gratitude rather than praise).

⁵³. Of course, on Wallace's account, there is an asymmetry between blame and praise such that believing an obligation has been violated renders appropriate the negative reactive attitude; morally worthy actions are not so intimately connected with the positive reactive attitudes (1994, 71–72). As Calhoun notes (2022, 13) the positive reactive attitudes are largely sidelined on Wallace's account.

(Section 1.3.3.1). What might praise seek recognition of, or shared valuing of, in these cases? Let's consider how one model that Calhoun critiques—Telech's—plays out in each case. Telech has it that praise is a moral invitation to co-value the action done, especially its meaning for the praiser; to accept credit for it, and feel directed pride. On this model, in praising, Gordon's colleague is inviting Gordon to co-value the bravery instantiated in the perceived overcoming of body shame; to accept credit for so doing, and to feel pride in so doing. Though perhaps not consciously intended, the praise also focuses on the social meaning that Gordon's supposed bravery has for the praiser. Gordon (2020) writes that:

Praising fat people's "bravery" and "confidence" is a subtle kind of othering, a reflection of the speaker's values, biases, and limited understanding of fat people's experiences.

What becomes visible in applying this model to a specific case of non-ideal praise is that the praiser is also falsely assuming that Gordon shares the presuppositions of the invitation to co-value. The praiser communicates the normative expectation that Gordon feels shame about her body, and presupposes that this shame is apt. The praiser's values—fear of fatness, assumptions about shame, and Gordon's relationship to her body—are expressed. Gordon cannot accept the praise without accepting these presuppositions, and endorsing these values.

Consider next the "daddy dividend" case. In praising, the stranger on the subway is inviting Lindauer to co-value his supposedly elective performance of parenting duties; to accept credit for doing so and feel pride. The praiser, again tacitly and perhaps unintentionally, is expressing something about the social meaning of the behavior for them: that it has exceeded their normative expectations about the gendered roles within families.

What becomes visible here is that the praiser is falsely assuming that Lindauer shares the presuppositions of the invitation to co-value. The praiser communicates the view that a father doing basic parenting is doing some elective good—something above and beyond what is normatively expected of him. The praiser's values—endorsing gendered normative expectations for parenting—are expressed. Lindauer cannot accept

the praise without accepting these presuppositions, and endorsing the values.⁵⁴ It is via these presuppositions that praise expresses certain values, and how, with uptake, these values are affirmed. The presuppositions about value expressed in praise, then, are one way in which social pressure to affirm certain values—in these instances, oppressive values—is transmitted.

The idea that praise involves normative pressure is not wholly absent from existing analyses. Telech focuses on the normative pressure that is present in praise's invitation—the pressure to accept the invitation unless there are good reasons not to do so. Quoting Martin (2019, 11), Telech points out that “valid invitations... ‘carry with them a certain legitimate pressure to accept, where the invitee needs a good reason to refuse (beyond say “I don’t feel like it”)” (2021, 171). But all the pressure, on this view, comes from the form of the invitation itself: the request which requires reasons to be given if refused (“Why don’t you want to co-value that good deed with me?”). Looking at non-ideal cases of praise—cases in which it operates within oppressive norms or relationships—we can see that there is another source of pressure in operation: namely, the pressure exerted by the presuppositions of the expressed praise, that the praised (and wider audiences) share the same view and values as the praiser. The expression of praise exerts social pressure to accept the presuppositions of the praise: that Gordon should be ashamed of her body, that fathers’ basic parenting is exceptional.

This form of pressure is rendered invisible on Telech’s analysis, which abstracts from the particular values operative, instead “[employing] as a placeholder the term ‘laudable standard’” (2021, 162).⁵⁵ By filling out the putative laudatory standards (or normative expectations or values) expressed in instances of praise as we find them in non-ideal circumstances—and in particular in instances in which the praisee does not endorse the values expressed, and does not share the presuppositions of the expressed praise—we see that praise serves an important function in our social practices. It functions to exert social pressure (albeit tacitly, by way of presupposition) towards accepting certain values and expectations. This resonates

^{54.} Another possible interpretation of this case, given what I have argued about the aptness of praise for merely accountability-responsible persons, is that the praiser is assuming that Lindauer is merely accountability-responsible, and is praising by way of helping to develop his capacities and sensibilities. The presupposition of this praise is quite different, but still problematic: that, qua father (and in contrast to those of mothers), Lindauer’s parenting capacities are naive and in development. Thus the praise has patronizing presuppositions. See Jeppsson and Brandenburg 2022 for discussion of patronizing praise.

^{55.} Note that Telech’s analysis focuses on the good will that is manifested in meeting the relevant laudatory standards (whatever they are) (2021, 161–62). This seems to me unduly narrow, since various instances of praise (and other reactive attitudes) focus on qualities of character displayed, rather than quality of will expressed.

with the idea, from Bennett Helm (2019, 185–86, 188), that positive reactive attitudes (Helm focuses on gratitude) serve a social function: affirming norms or expectations, and entrenching motivations to act in accordance with certain values. But since those values and expectations may sometimes embody oppressive values and norms, praise may function to entrench oppressive values.⁵⁶

4.4.2.2 Signaling

A second mechanism through which reactive attitudes might function to exert social pressure is by signaling. This requires expanding the analysis of praise to consider what it communicates to third parties: doing so, again, embraces Calhoun’s approach, whereby we consider the social practices as we find them. But it helps us (contra Calhoun) see the point of the uptake sought by the positive reactive attitudes. The focus of the analyses of praise with which Calhoun engages is communicative praise: praise addressed to another where uptake is sought. The form of communicative act at issue, in Calhoun’s discussion, and that of the authors she engages with, however, is second-personal in form: the praiser is addressing, and seeking uptake from, the putatively praiseworthy person. As such, the function of praise, and what it signals, for third parties is not within the purview of these analyses; Telech writes that:

To the extent that [an instance of] recognition does not seek uptake from the praiseworthy agents, it will not be an instance of praise in the moral address sense of interest to the communicative theorist. (2021, 166)

What might this focus on second-personal communication omit? In a recent functionalist account of blame, Shoemaker and Vargas have argued that the disparate cases and kinds of blame are best unified by understanding blame as a kind of signal. On their view, the function of blame (in its various forms) is to signal the agent’s competence and commitment: competence, qua moral agent, in understanding moral norms, recognizing violations of the relevant norms, and in policing those norms. And expressing blame signals the agent’s commitment to certain moral norms, and to enforcing them. Crucially, any one expression of blame might signal multiple

⁵⁶ One further way in which expressions of praise may exert pressure, articulated by Grahle, is by providing “petitionary [in the sense of ‘approaching someone with a concern’] reasons” (2019, 155). The idea is that second-personal expressions of admiration for some feature provide *pro tanto* reasons for maintaining or developing that feature. As Grahle notes, things can go wrong when the feature towards which admiration is expressed is not really admirable, or when expressions are shaped by sexist or racist norms. Then expressions of admiration might provide reasons to maintain features that ought not to be maintained or developed (2019, 161).

things to multiple audiences: blame to the direct audience;⁵⁷ solidarity to the victim, a warning to would-be norm violators, affirmation of the norm and community with those who endorse it (2019, 10).

Since praise is part of the same moral ecosystem as blame, we can consider what praise might signal. Abstractly, we might suppose that expressions of praise signal the praiser's competence in identifying when certain norms or values have been promoted, and their commitment to—and affirmation of—those values or norms (in particular a commitment to promoting them, and celebrating other's promotion of them). Less abstractly, let us again consider what is signaled in the two cases of praise introduced above. Where Gordon's colleague praises Gordon's putative bravery, the praiser signals her commitment to certain values and norms: norms according to which bravery is valued, and the promotion of such values. But also, as we have seen, the praiser will signal commitment to the presumptive evaluative framework according to which fat bodies carry shame that requires bravery to overcome. Hence Gordon's discomfort with—and ultimate rejection of—the praise: accepting it would signal commitment to values that she rejects, and indeed, values which denigrate and degrade her.⁵⁸

In the “daddy dividend” case, the praiser signals their commitment to celebrating men who do more parenting than, according to the praiser's evaluative framework, is required. But this also signals their commitment to that evaluative framework, with its gendered presuppositions about parenting roles. Hence Lindauer and Khader's critique of the praise: accepting it would signal endorsement of, and commitment to, the gendered evaluative presuppositions of the praise.

How do these remarks about the signaling function of praise relate to Calhoun's project? The point is that, in identifying the signaling role of praise, we can see why there is in fact purpose in drawing attention to the value promoted, and seeking uptake both on the part of the person directly praised, and in third parties to the praise, in recognizing and emotionally responding to those values. In doing so, praise is fulfilling a function of signaling commitments to certain values and norms, and seeking to affirm the presupposed evaluative frameworks expressed in that praise. Through such signals and affirmations, oppressive norms can be entrenched, albeit unintentionally. We can make sense of why praise might seek to affirm values—

^{57.} Though on their view there need not be any dyadic communicative relationship between the blamer and the blamed, as when we blame the dead, or blame absent parties. One can also blame oneself, or engage in private blame.

^{58.} As mentioned in Section 1.3.2, there may be other scenarios in which such praise signals a shared rejection of oppressive norms.

oppressive or otherwise—both for the person praised, but also with third parties to the praising interaction, as well as for oneself.

Praise may, as Calhoun argues, respond to and embody a conception of persons as capable of taking responsibility. It may, as Eshleman argues, sometimes respond to a conception of responsible persons that embodies the “‘extra’ capacity agents might enjoy above the minimum” (Eshleman 2014, 236). But whether praise is responding to accountability-responsible persons—sensitizing merely developing capacities—or responsibility-takers, in both cases it plausibly has a role in exerting social pressure, both in its presuppositions, and in signaling to others. And given this, it makes sense that it seeks uptake in recognitive and emotional ways. The praiser cannot, or should not, assume that the target shares the values expressed. This social pressure may in many cases be valuable, affirming and deepening our (the praiser, praised, and third parties’) commitments to shared values. But the social pressure, as in the cases we have examined, may function in disvaluable ways—presupposing and attempting to affirm oppressive values and norms.

4.5 Social Practices

In this commentary, I have embraced many aspects of Calhoun’s endeavors to expand the conception of the practices we consider to be responsibility practices, and the (models of the) responsible persons and responsibility-recognizing attitudes that populate these practices. However, I have argued that if we are to start with the social practices, broadly construed, then we have good reason to start with the non-ideal forms of social practices we find, and the ways in which social hierarchy and power dynamics inflect our responsibility practices.

Doing so led me to make the following observations about Calhoun’s tripartite model of responsible persons. First, the claims Calhoun makes about the default status as accountability-responsible, and default attitudes of trust in compliance responsibility, seem not to be borne out under oppressive social contexts (Section 1.3.1). Second, oppressive norms can distort the responsibility-recognizing attitudes (reactive ones such as praise and blame, as well as prospective attitudes towards responsibility-takers) turning them into tools of oppression that threaten the value of the status of responsible person. Moreover, in the context of oppressive norms, the alignment between the responsibility-recognizing attitudes and the tripartite model of responsible persons breaks down (Sections 1.3.2, 1.3.3). So, whilst Calhoun draws attention to a number of aspects of our practice that may have value, it is important to attend carefully to the contexts in which those values can be realized, or distorted.

Even in non-oppressive contexts, I argued, the welcome expansion of accountability responsibility to include those with merely developable capacities should lead us to endorse the idea that positive responsibility-recognizing attitudes, such as praise, are apt not only for responsibility-takers, but also for those with developable capacities (section 1.4.1). Moreover, a feature of praise most visible where reactive attitudes are used to enforce oppressive values, but also present in other cases, is its role in exerting social pressure—in “affirming and deepening our commitments” as some authors have put it. I teased out the mechanisms by which praise might do this—via the presuppositions of expressions of praise and via signaling commitments to certain values. If praise has these functions, then even if it does not only embody a conception of the responsible person as accountability-responsible, we can make sense of the importance of praise seeking recognitive and emotional uptake (Section 1.4.2), both in second- and third-personal interactions. This will be particularly important where praise is a form of feedback for accountability-responsible persons with developing moral capacities.

I think it is striking that analyses of praise (and the reactive attitudes in general) have given little attention to the role of reactive attitudes as norm entrenchment and enforcement mechanisms.⁵⁹ Perhaps one reason for this is the tendency to move away from understandings of our responses to good and bad doings in crudely consequentialist terms, whereby praise and blame are seen as carrot and stick, reward and sanction.⁶⁰ Instead, the exalted features of responsible persons—as sensitizable to norms, as literate in nuances of normative contexts, persons who elect to promote values—are given prominence—and not without good reason. However, I hope that attention to the social practices more widely, and in particular the social practices in their non-idealized, and often oppressive manifestations, shows that it is worth attending to the more mundane parts of our responsibility practices: the parts that function as the levers of social pressure, entrenching values and norms—for better or for worse.

In relation to Calhoun’s project of better articulating what is valuable, via an expanded conception of our responsibility practices and of responsible persons, I hope to have shown that a less idealized starting point can help us to articulate a more equivocal view of the value of those aspects to which Calhoun draws attention.

^{59.} Though in fact Calhoun’s (1989) distinction between justification and point, in the context of moral responsibility ascriptions and responses, gets just at this social pressure, I think: that there might sometimes be a point in reproaching (moral education, in entrenching and motivating adherence to norms), independently of the justification for doing so (405–406).

^{60.} Indeed, I think there are good reasons to move away from such crudely construed models, as I argue in Holroyd 2007.


This enables us to engage with the question of which social contexts better realize these potential values, and which undermine them. Attention to our non-idealized social practices can inform our articulation of these practices, and the value of the status of responsible person; and, where needed, improve both the conceptualizations of them, and in turn the practices themselves.

4.6 References

- Anderson, Elizabeth S. 2010. *The Imperative of Integration*. Princeton: Princeton University Press.
- Calhoun, Cheshire. 1989. "Responsibility and Reproach." *Ethics* 99 (2): 389–406.
- . 2022. "Appreciating Responsible Persons." In *Oxford Studies in Normative Ethics*, vol. 11, edited by Mark Timmons, 9–28. New York: Oxford University Press.
- Campbell, Ellie. 2022. "I Made a Name for Myself with 'Sex-Positive' Comedy. Then I Was Raped on a Night Out. Would My Openness Be Used against Me?" *The Guardian*, August 6, 2022. <https://www.theguardian.com/society/2022/aug/06/sex-positive-comedy-las-vegas-rape>.
- Carbonell, Vanessa. 2019. "Social Constraints on Moral Address," *Philosophy and Phenomenological Research* 98 (1): 167–89.
- Ciurria, Mich. 2020. "The Mysterious Case of the Missing Perpetrators: How the Privileged Escape Blame and Accountability." *Feminist Philosophy Quarterly* 6 (2).
- . 2021. "A New Ameliorative Approach to Moral Responsibility." *Verifiche: Rivista Trimestrale di Scienze Umane*, 1 (2): 159–82.
- Coates, Justin. 2019. "Gratitude and Resentment: Some Asymmetries." In *The Moral Psychology of Gratitude*, edited by Robert Roberts and Daniel Telech, 160–75. London: Rowman & Littlefield.
- Corpus, Jennifer Henderlong and Kayla A. Good. 2020. "The Effects of Praise on Children's Intrinsic Motivation Revisited." In *Psychological Perspectives on Praise*, edited by Eddie Brummelman, 39–46. Abingdon: Routledge.
- Eshleman, Andrew S. 2013. "Worthy of Praise: Responsibility and Better-than-Minimally-Decent Agency." In *Oxford Studies in Agency and Responsibility*, vol. 2, edited by David Shoemaker and Neal A. Tognazzini, 216–42. Oxford: Oxford University Press.
- Eskreis-Winkler, Lauren and Ayelet Fishbach. 2020. "When Praise—versus Criticism—Motivates Goal Pursuit." In *Psychological Perspectives on Praise*, edited by Eddie Brummelman, 47–54. Abingdon: Routledge.
- Glasgow, Joshua. 2007. "Three Things Realist Constructionism about Race—or Anything Else—Can Do." *Journal of Social Philosophy* 38 (4): 554–68.
- Gordon, Aubrey. 2020. "Please Don't Call Fat People 'Brave' Just for Existing." *Self*, October 12, 2020. <https://www.self.com/story/fatness-bravery-bias>.
- Grahle, André. 2019. "Admiration and Normative Support." In *The Moral Psychology of Admiration*, edited by Alfred Archer and André Grahle, 149–63. London; New York: Rowman & Littlefield.
- Haslanger, Sally. 2015. "Social Structure, Narrative, and Explanation." *Canadian Journal of Philosophy* 45 (1): 1–5.
- Helm, Bennett W. 2019. "Gratitude and Norms: On the Social Function of Gratitude." In *The Moral Psychology of Gratitude*, edited by Robert Roberts and Daniel Telech, 176–94. London: Rowman

& Littlefield.

- Holroyd, Jules. 2007. "A Communicative Conception of Moral Appraisal." *Ethical Theory and Moral Practice* 10 (3): 267–78.
- . 2018. "Two Ways of Socialising Responsibility: Circumstantialist and Scaffolded-Responsiveness." In *Social Dimensions of Moral Responsibility*, edited by Katrina Hutchison, Catriona Mackenzie, and Marina Oshana, 137–62. New York: Oxford University Press
- . 2021. "Oppressive Praise." *Feminist Philosophical Quarterly* 7 (4).
- Jeppsson, Sofia and Daphne Brandenburg. 2022. "Patronizing Praise" *Journal of Ethics* 26: 663–82.
- Khader, Serene and Matthew Lindauer. 2020. "The 'Daddy Dividend': The Gender Division of Labour and Regression Towards Patriarchy." *APA Newsletter on Feminism and Philosophy* 19 (2): 6–8.
- Manne, Kate. 2020. *Entitled: How Male Privilege Hurts Women*. New York: Crown Publishing Group.
- Macnamara, Coleen. 2013. "'Screw You!' & 'Thank You.'" *Philosophical Studies* 165 (3): 893–914.
- McGeer, Victoria. 2019. "Scaffolding Agency: A Proleptic Account of the Reactive Attitudes." *European Journal of Philosophy* 27 (2): 301–23.
- Roberts, Dorothy. 1997. *Killing the Black Body: Race, Reproduction, and the Meaning of Liberty*. New York: Pantheon Books.
- . 2022. *Torn Apart: How the Child Welfare System Destroys Black Families—And How Abolition Can Build a Safer World*. New York: Basic Books.
- Shoemaker, David and Manuel Vargas. 2021. "Moral Torch Fishing: A Signaling Theory of Blame." *Noûs* 55 (3): 581–602.
- Soenens, Bart and Maarten Vansteenkiste. 2020. "Understanding the Complexity of Praise through the Lens of Self-Determination Theory." In *Psychological Perspectives on Praise*, edited by Eddie Brummelman, 27–35. Abingdon: Routledge.
- Strawson, P. F. 1962. "Freedom and Resentment." *Proceedings of the British Academy* 48: 1–25.
- Telech, Daniel. 2021. "Praise as Moral Address." In *Oxford Studies in Agency and Responsibility*, vol. 7, edited by David Shoemaker, 154–81. Oxford: Oxford University Press.
- Titheradge, Noel. 2022. "Scores of Safeguarding Complaints at After-School Clubs." *BBC News*, 5 May, 2022. <https://www.bbc.co.uk/news/uk-61325477>.
- Today in Focus. 2022. "Why Comedian Grace Campbell Refuses to Be Silenced about Sex." *The Guardian*, 19 August, 2022. Podcast, 24:31. <https://www.theguardian.com/news/audio/2022/aug/19/why-comedian-grace-campbell-refuses-to-be-silenced-about-sex-podcast>.
- Vargas, Manuel. 2013. *Building Better Beings: A Theory of Moral Responsibility*. Oxford: Oxford University Press.
- Wallace, R. Jay. 1994. *Responsibility and the Moral Sentiments*. Cambridge, MA; London: Harvard University Press.



5

BEING PREDICTABLE, BEING TRUSTWORTHY

Heidi L. Maibom

Department of Philosophy, University of Cincinnati and ILCLI &
Department of Philosophy, University of the Basque Country

It is my great pleasure to comment on Cheshire Calhoun's *Descartes Lectures*. I have long admired her work and am especially sympathetic to her approach to morality, which focuses on the community over the lone search for the moral law within (see e.g. Calhoun 2004; Maibom 2010). In this paper, I focus on her second lecture, which is concerned with spelling out a notion of compliance responsibility that goes beyond praise and blame. Being compliance responsible is not simply a matter of having the capacity to discern, learn, and enact what is expected of us, but also of *exercising* that capacity, Calhoun says. Compliance responsibility enables social life because it gives people reason to suppose that others will, as a matter of fact, adhere to the common norms. They can then simply trust that this is so and go about their normal lives without having to consider, specifically, whether this or that person will actually act thus-and-so. Quoting Calhoun directly:

The default presumption in everyday social life is not just that those we interact with can learn and can comply with normative expectations and thus are liable to blame once they are familiar with what is normatively expected. The default presumption is, in addition, that they have already learned what's normatively expected and are disposed to comply with those expectations—at least the most basic ones. Thus, the *expectations* to which we are prepared to hold people to account are not just normative expectations about how people ought to behave. They are also empirical, predictive expectations about how people largely will behave. (30)

In what follows, I will show how we can get to some of the same conclusions Calhoun reaches coming from a different direction, namely the research on folk psychology within the philosophy of mind. This helps strengthen Calhoun's case, and it shows how greater integration of research within philosophical sub-areas helps move the debate forward. I also have some reservations or, perhaps, additions to Calhoun's argument.

5.1 What Allows Social Cooperation and Coordination? The Sociopolitical Myth

What allows social cooperation and coordination is an issue much researched in anthropology, psychology, philosophy, biology, mathematics, decision theory, and political science (Axelrod 1984; Binmore 1994; de Waal 1982; Hrdy 2009; Richerson and Boyd 2005; Sterelny 2012). One might be interested in how cooperation arose in our human ancestors (how it evolved), in how it *could* evolve among rational participants, or how it *actually* works on the ground. One might want to provide a proximal explanation, or, alternatively, a distal one. After all, it is quite possible that the reason cooperation evolved has relatively little to do with why we are still, as a matter of psychological fact, cooperating today.

It is common to think of the origin of cooperation as intimately related to morals and politics. Thomas Hobbes famously insisted that the state of nature—the state anterior to elaborate systems of cooperation enshrined in a moral or political structure—was characterized by a “war of every man against every man” ([1668] 1994, XIII.13). Life in the state of nature, therefore, was “solitary, poor, nasty, brutish, and short” ([1668] 1994, XIII. 9). Hobbes was well aware that he was describing a rather idealized affair, but nonetheless thought it was a good way of understanding the state. It is a brilliant way of justifying why people should transfer their self-sovereignty onto a higher authority. Because, if the alternative is as Hobbes describes it, then the restriction of liberty that inevitably ensues upon entering into a social system is almost certainly worth it.

We can equally apply a Hobbesian way of thinking to any moral system. Morality, as we know, binds and restricts, but it also safeguards (see also Freud [1930] 1961). Overall, however, the benefits are supposed to outweigh the costs. Whether that is actually true is an empirical question rarely considered in mainstream philosophy. A recent book by David Graeber and David Wengrow (2021) argues that some moral systems are much more disadvantageous to most individuals than the presumed state of nature of so-called “savages” (they use First Nations tribes from Eastern Canada at the time of the first European settlements as an example). Of course, we now

know that even “savages” have their own moral systems, which are sometimes less restrictive and kinder than their European equivalents, unless, of course, you happen to be a woman (Graeber and Wengrow 2021).

A cursory look at the moral and legal norms that prevailed in Europe during most of its history is unlikely to persuade one that the majority of people were better off with them than without them. Slavery was common across Europe for thousands of years and, after that, feudalism. Norms concerning property transfer, ownership, and inheritance primarily served to ensure the continued wealth and power of the few, and to ensure a permanent source of cheap employment (sounds familiar?) and the subjugation of women. Women have been barred from getting an education, inheriting, owning property, voting, being part of the clergy, choosing their husbands, and the right to raise their children (see e.g. de Beauvoir [1949] 2011). They have been beaten and murdered, often with the backing of prevailing socio-moral norms. I feel a rant coming on, but I shall try to contain myself. The greater point I’m trying to make is that norms are tricky things. They are without doubt useful in social coordination, but operate just as well as tools of oppression and suffering.

A common way to think about norms is, as Hobbes suggested, that they are essential for ensuring social cooperation because each individual is only concerned with their own good and, possibly, the good of their family (see e.g. Axelrod 1984; Hobbes [1668] 1994). In the final analysis, this leads to a highly unstable system, which makes life for everybody harder and more unpleasant than it needs to be. A system of norms imposes upon individuals *some* concern for others and the common good. So, the real reason we need norms is because humans are naturally too selfish to cooperate in the absence of a social structure that enforces norms of cooperation. These norms, then, are for the good of everybody. The story can vary in details, but this is the essence of it. I have given some reasons above, from the *actual* history of morality, to suggest that a number of social systems systematically *disadvantage* many, and perhaps even most, people, and women in particular. So, the happy story of how we came from the brute state of nature into the fold of morality may be little more than fancy advertising for our current socio-political system.

This skepticism about the goodness of morality is the first thing I would like to bring to the attention of Calhoun, for she paints too rosy-red a picture of norms, I think. Social practices, she argues, are structured by norms. In some cases, such norms are constitutive of such a practice, meaning that if the norms were not followed, the practice would cease to exist. In other cases, they are matters of common decency. And although norms may not all be good, the vast majority of examples she provides

are norms her readers would readily accept as *morally good*. Moreover, when she goes on to talk about the fact that our normative expectations are mostly *default*, she references Baier's work that "reminds of how pervasive reliance on others' actual, even if possibly minimal, good will in fact is" (39). And, although she allows for the fact that not all norms are "*legitimate*" (46), the underlying assumption seems to be that the force of (at least our) moral norms is generally a good thing for everyone involved. But I think it is important to temper such optimism by a sober look at how norms *actually* operate.

Right now, in the U.S., conservative forces are systematically chipping away at freedom of speech and expression (as long as it doesn't involve companies buying influence with the government, or the carrying of arms). Books are being banned (Ciabattari 2023), even fantasy novels such as Philip Pullman's *Northern Lights* (apparently for its atheist tone), state governments are attempting to interfere with state university curricula and forbidding teaching race or gender (Staver 2021), women are being forced to carry pregnancy to term with few, if any, exceptions (Bellware 2022; Watkins 2023), and others are having forced upon them a sexual identity alien to them (Yurcaba 2023). On the left, any questioning of issues related to gender, race, or sexuality that challenges the standard position, are viciously countered and the questioners "cancelled" (Weaver 2023). None of this is done in the name of evil or oppression, but in the name of *Morality*. We can comfort ourselves that *we* are talking about *the real* morality, but it is becoming increasingly difficult to see how this is useful.

Moreover, one might ask oneself why we should be so delighted with cooperation in general. The idea, we are told, is that when we cooperate, we all benefit. But this evidently depends greatly on the particular system of cooperation in place. Frequently, systems of norm-governed cooperation favor the powerful over the less powerful, even if they offer some protection to the latter. And, once in operation, these norms cannot be opted out of. The norms that have traditionally governed a woman's access to resources have hardly been to her advantage, such as being barred from having an education, a profession, a divorce, or what have you (see e.g. de Beauvoir [1949] 2011). It was hardly to the advantage of the common man that he could vote only when in possession of land or wealth. The Occupy Wall Street movement voiced Americans' concern that the socio-political *status quo* caters to the 1 percent and systematically oppresses the 99 (NPR Staff and Wires 2011; Anthony 2021). The movement was broken up by the coercive powers of the state (a.k.a. the police). Where am I going with this? I think there is an excellent case to be made that morality is a system of social control which benefits those who currently

wield power (cf. Nietzsche [1887] 1989). And it remains an open question whether such a system benefits more people than it disadvantages. Calhoun's argument does not hinge on systems of social norms benefitting the many, but if they do not, it rather does affect the way we regard her proposal. We might, for instance, regard compliance responsibility as a liability and a virtue in equal measure.

5.2 Social Cooperation and Coordination: The Philosophical Myth

We can, however, come at the “problem” of human cooperation from an entirely different direction, which I now will go on to describe. I do so because I think it provides a different outlook on *why* we have norms—and in particular such a numerous and varied collection—and because it might encourage new ideas about the connection between responsibility and norms. The direction I have in mind is that of folk psychology. In the 1980s and early 1990s, the only game in town—at least according to Jerry Fodor (1989)—was the view that folk psychological theory allows us to predict the behaviors of others. Without it we would be at sea in the social world. Indeed, society as we know it would collapse. Why, you might ask. Good question. Philosophers have long assumed that non-human animals, if not simply foaming at the mouth and driven by blind instinct, are certainly incapable of anything remotely resembling sophisticated social behavior. Yet, many of those animals went on living with each other in groups and coordinating their behaviors, avoiding capture, stalking, mating, raising viable offspring, and some even living with humans in productive relationships, academic skepticism notwithstanding. All this would suggest that at least basic forms of interaction are possible without anything as complex and sophisticated as a theory of folk psychology. At the end of the last century, however, animal social organization, if it was thought to exist at all, was deemed to be almost entirely irrelevant to human social behavior and cooperation.

Okay, so why do we need to predict others' thoughts, feelings, intentions, desires, or actions in order to cooperate with them? Two answers immediately present themselves. We need such knowledge in order to *coordinate* behavior and in order to avoid being exploited. On the first point, we can note that although coordinated behavior is widely on display in non-human animals, human social coordination is vastly more complex, sophisticated, and deadly. So, there is certainly an argument why the forms of cooperation we see in humans are special and require a distinct form of explanation compared to those we might deploy to explain cooperation in the most sophisticated non-human animals. Nonetheless, the standard claim is that *all* human cooperation must be explained this way, something which is highly doubtful given what we know of non-human animal cooperation. But let's be

charitable and assume that what these philosophers had in mind were complex forms of social coordination.

Suppose, then, that whenever you are about to do anything “sophisticated” with someone else, you first predict what they will do, so that you can conform your behavior accordingly. However, this is easier said than done because, by the reasoning above, your potential cooperator will also base his action on what he thinks you will do. But, since you don’t know what you are going to do, he cannot know either. This implies that you cannot base your decision on his decision either, because it has not yet been made. Weirdly, then, you can only predict what the other person will do *once* you have decided what to do, which defeats the very purpose of the project of “reading” the other person’s mind! At least on this picture. This is the cooperation problem, brilliantly exposed by Adam Morton (1996).

Morton was particularly concerned with the modeling of social cooperation that arises from the prisoner’s dilemma. The prisoner’s dilemma models the risks and rewards of cooperation. In such a dilemma, each person must decide whether or not to cooperate with the other without any certainty that the other person will cooperate in turn. The canonical story goes something like this. Two suspected criminals who have jointly committed some crime, say a bank robbery, are questioned by the police. To prevent them from collaborating on some story, they are each in their own room (and therefore unable to communicate with one another). In order to cajole a confession out of them, the police give them the following options. If prisoner A confesses but B does not, prisoner A will go free. But in this case, B will go to jail for 8 years. And *vice versa*, of course, if prisoner B confesses but A does not. If they refuse to confess, they both receive a term of, say, 2 years each (note: do not ask a legal scholar to evaluate this example!). If both confess, then each will receive a term of 5 years. Whereas the optimal *mutual* strategy would be for the prisoners *not* to confess, self-interest demands that each should safeguard against receiving the so-called sucker’s payoff, i.e., 8 years in jail. The only way to do that, however, is to defect. The best-case scenario, then, would be one in which *the other prisoner* confesses, but you do not.

However, it might seem that if one could figure out what the other person is going to do in a prisoner-style dilemma, one can make a better decision. It is here that the ability to predict what others will do, or decide to do, ought to figure. The incentives being what they are, the rational prisoner will only not confess if she believes that the other person will not confess. Folk psychology ought to help her with this. But Morton points out that this is actually not a solution at all. For to figure out

what the other person will do, she must consider what it is rational for the other person to do. And when she does so, she finds that this person equally will only not confess if she believes *she* will not confess. The result is a stalemate. The problem with the prisoner's dilemma is that the decisions of the players are fundamentally entangled. Neither person is able to reach a good decision about whether to defect or to cooperate in the absence of the other deciding to cooperate or defect.

It is widely observed that cooperation is a better strategy in *iterated* prisoner's dilemmas. Here the players can base their subsequent decisions on the previous moves of the other because they play an open-ended number of rounds. They can continue to cooperate (thereby rewarding the other for good behavior) or defect (punishing bad behavior), or change their strategy altogether. Instead of focusing on crime and punishment, iterated prisoner's dilemmas typically concern simple rewards, and the trick is to get as much of it as possible. Nonetheless, these dilemmas mirror the reward structure of the original one-shot prisoner's dilemma. It is to the advantage of both to cooperate, but the danger of being exploited by the other, so that one receives little or no reward, remains real. Typically, the reward structure in such scenarios makes cooperation rational due to the opportunity cost in defecting. Plausibly, if you defect, so will your partner in the next round, leading to reduced rewards overall for both of you (if you defect too, neither one of you will get much, but if you cooperate, you get the sucker's payoff). Moreover, initial defection on the part of one partner can easily spark a long row of defection on either side (one partner responding to the other partner's defection in the last move), leading to reduced rewards for both parties.⁶¹ In this kind of environment, it might be thought that the ability to predict the other person's upcoming move would be of great use. Why? Because if I cooperate blindly without any attention to what you do or intend to do, I might end up badly exploited. It only makes sense for me to cooperate if you, too, will cooperate. So, suppose I cooperate because it is in my interests to cooperate, and because it is in *your* interests to cooperate too. On that basis, I can predict that you will cooperate, as indeed you can predict that I will cooperate, and so we both cooperate. The problem, however, is that people are not perfectly rational. We see evidence of this all over the place. So, we need an extra step, namely that I *believe* that it is in my self-interest to cooperate and that you *believe* that it is your self-interests to cooperate too. This belief will presumably cause a desire to cooperate on both our parts, leading to an intention that can then play the appropriate role in prediction. Does this avoid the unhappy entangled decision-making we faced above? It doesn't seem to. Why? Because I form the intention to cooperate *only if* I predict

⁶¹. I describe the typical theoretical construal here. Whether that is actually matched by real-life behavior is not at all clear; cf. Raihani and Bshary 2011.

that you intend to cooperate, and you form the intention to cooperate *only if* you predict that I will cooperate. We have returned to the situation we were in before!

What is needed is a strategy for cooperation *that does not rely on predicting* the other person's intentions or actions. As Morton puts it, one ought to choose “an action, which is part of a combination in which no one can do better by unilaterally choosing differently” (2003, 5). And this is, in fact, exactly what standard decision theory gives you. Most people are familiar with *tit-for-tat*. Here your first move is always to cooperate. In all subsequent moves, you mirror the last move of your partner. According to Robert Axelrod (1984), this is the most successful and stable approach to such the iterated prisoner's dilemma. It has a number of features which makes it successful. It is nice (cooperation first), but willing to retaliate (*tit-for-tat*), and it is forgiving (if a partner starts cooperating, even after a defection, it will cooperate too).

Tit-for-two-tats can also give rise to stable and mutually beneficial patterns of interactions *without* having to fall back on mental state ascriptions. It is a reactive approach if you like. You simply decide to cooperate without worrying (overly) about whether the other person will cooperate or not, then see what the other person does. And you give them the chance to cooperate again after an initial defection. In this way, you don't get stuck being unable to decide whether to cooperate or not. In other words, you take a chance, twice, then adapt your behavior to that of your interaction partner. So much, then, for the assumption that we *need* to be able to ascribe psychological states to others to cooperate with them wisely. This, of course, does not rule out that some amount of theory of mind is needed to figure out what particular form of cooperation is at issue but, again, this really isn't a matter of *prediction*. Why? Your decision to cooperate is not based on a prediction of what the other person will do—because, as I noted above, your decisions are entangled—but on adopting a certain approach to a cooperative problem, such as mirroring your partner's previous move. It seems likely that this type of investment is contingent on the partner's investment. When carefully interpreted, there seem to be many cases in which animal cooperation is best explained in terms of *tit-for-tat* strategies, which require some amount of vigilance (to detect defection, for example) (Raihani and Bshary 2011).

Now, in large-scale societies, people often interact with others with whom they are unlikely to have further interactions. Such interactions are, then, better modeled on a one-shot prisoner's dilemma, rather than an iterated one, raising questions about how cooperation can get off the ground in the first place (Frank 1988). One

answer might be that the things that get cooperation going in the first place, such as relatedness or attachment, are not well modeled on a prisoner's dilemma (Raihani and Bshary 2011). Once you *do* have some cooperation, it may be easier to build on it, creating cooperative practices on a larger scale that are not narrowly tied to logic, but possibly also sentiment (Frank 1988) or morality. In these cases, one might feel guilt at the prospect of not tipping the waiter in a restaurant one will never return to (Frank's example; Frank 1988) or think that it is not "the right thing to do." If so, one's decision to cooperate is not based on a prediction of what the other person intends to do. As such, they represent solutions to the intractable problem of entangled mutual decisions.

Predicting someone's intention or action, then, is not only not required to solve classic cooperation problems, but is also often impossible given the entangled nature of cooperation problems. Now it might be thought that, when you engage in a tit-for-tat strategy, you are as a matter of fact *predicting* that the other person will cooperate or defect in their next move based on the previous move. The tit-for-tat strategy does not *require* such prediction, of course. Interestingly, much animal cooperative behavior can be understood on a tit-for-tat model, but is generally *not* explained in terms of psychological prediction or other involved cognitive mechanisms (Raihani and Bshary 2011). Since we quite likely evolved after some kinds of cooperative behaviors—at the very least involving social coordination—evolved in our ancestors, there would be no need to overlay the relevant strategies with costly cognitive machinations. Moreover, momentary reflection on why people stop cooperating in response to non-cooperative partners shows that this is often based on anger, outrage, or vengefulness, and has relatively little to do with the prospect of future defection on behalf of the other person. None of this shows that prediction can play *no* role, of course, but it does rather indicate that it plays a much smaller and peripheral role than it has made been made out to.

Calhoun seems to intuit this reduced role of prediction when she argues that the better explanation for someone choosing to cooperate is that they have been taught to do so. That is, it is unlikely that people are taught to cooperate *as such*. Instead, they are taught, or just learn by observing others, to pay at the cash till, to send merchandise upon receipt of payment, to signal right when turning right, to accept the authority of the policeman to tell them what to do, to bring a bottle of wine when invited for dinner, and so on. Moreover, it is not just something they have been taught so that they can do it, or not, should the situation arise. It is, as Calhoun rightly points out, a norm. It is *what* is to be done or, alternatively, *how* it is to be done. Some of these norms are backed by legal sanctions. But even when they are

not, not living up to them comes at a cost: to one's reputation or to one's future opportunities to cooperate. Social sanctions are by no means negligible. But, as Calhoun also maintains, people don't simply follow norms because they fear sanctions. That type of motivation constitutes, at best, an early developmental stage (Kohlberg 1984). Instead, doing the right or correct thing is enough for most people. It is what you buy into, as it were, once you engage in certain transactions. If, for instance, I am an art dealer, I show art and sell it to people for profit. I share that profit according to an agreement with the artist. I take the money first, then send the art work to the buyer (unless they can take it with them). This is *what it is* to be an art dealer. So, all sorts of social structures support norm adherence. Because George takes me to be a customer and I him to be an art dealer, the space of our interaction is restrained by what it is to be one or the other. I can, you might say, *rely* on George to send me the picture I paid for, just as he can rely on my payment going through. Neither George nor I rely on one another because we have squared each other up and decided that we are each reliable. Instead, we rely on the larger social system within which our particular interaction takes place. In a way, our individual agencies have been absorbed by the larger system, and we each do what the respective roles we occupy mandate. Here there is no need to predict *the mental states* of the other (although you are certainly allowed to make predictions, should you want to do so).

This alternative way of thinking about what we need to cooperate has also been mirrored in new ways of thinking about folk psychology. Whereas the ascription of propositional attitudes, especially beliefs and desires, was all the rage in the 1980s and 1990s, there is now talk of mental models of various sorts helping us do the epistemic work of understanding others (Maibom 2003; 2007; Godfrey-Smith 2004; Spaulding 2018). For instance, I have argued that behavioral models help us extract meaning from the movement of living creatures, and that social models play an important role in allowing us to interact appropriately with other people within social contexts (Maibom 2007). Social structures take away the need to ascribe psychological states to others for predictive purposes. One simply identifies the right model. You are in Florence at a busy coffee bar. What do you do? You first pay the cashier, who gives you a receipt, which you subsequently present to the barista as proof of payment. Then you tell him exactly what you want. You can then relax (although this step is optional) and wait to be served some of the best coffee you will get in Europe. The barista might strike you as surly and unenthusiastic about his job (no West Coast tattooed hipster baristas here), but that does not affect the nature of the interaction. You pay and get a great coffee either way. Your barista is performing the role of barista (in Italy) and you perform the role of customer.

Once we shift the discussion from beliefs, desires, and intentions to social structures and circumscribed roles and interactions, we are changing the nature of the debate. And we are inching closer to the proposal that it is social norms, more than anything else, that makes the human world go around (i.e., that allows the seamless social interaction we observe in large-scale industrial societies). In fact, some have suggested that folk psychological ascriptions are less for prediction or explanation and more for social control. Or, to put it nicely, it is about “mind shaping” more than “mind reading.” Tad Zawidzki (2008), for example, suggests we “offload” the intractable epistemic task of deciding whether to cooperate or not onto the environment. By talking about others the way we do, we signal what behavior is acceptable, what one ought to think and like, and so on. Instead of resigning ourselves to the intractable task of trying to predict someone’s actions on the basis of their psychological states, often *ex novo*, we set up an environment that invites certain kinds of behaviors and discourages others. We can then rely on norm adherence *when it comes to social coordination and cooperation* instead of psychological prediction. A different way of capturing the basic idea is that, instead of honing our own skills at psychological prediction, we create predictable people by means of the incentives and disincentives built into the environment.

Tori McGeer (2007) goes even further and suggests that folk psychology plays a regulative role in helping conform people’s behaviors to reason, or a certain view of reason, and norms. “Our ways of organizing our environment, our ways of conducting ourselves in special orientation to one another, our way of using voice and body, our ways of dressing, all come normatively guided,” she says (147). But the action goes both ways. It is not just that we are trying to make other people conform to norms. She writes, “Skilled folk psychologists are not just able to read other people in accord with shared norms; they also work to make themselves readable in accord with those same norms” (148). Morton (2003) suggested something similar in his *The Importance of Being Understood: Folk Psychology as Ethics*. Why should we make ourselves “readable” to others? So that they can understand whether or not we are reliable interaction partners. By reducing the interpretive work that others have to do, we make ourselves more comprehensible to them. If we do not, others may deem us to “fall outside the realm of subjects [they] can interact with as free and responsible agents, able to make commitments to [them] or to understand the commitments [they] make to them” (McGeer 2007, 149). The background idea, of course, is that social life requires extensive cooperation with others, and so our motivation to be deemed decent partners in such cooperation is great.

What characterizes these suggestions is a general agreement that social life simply wouldn't be possible if there were no norms, solutions with norm-like structure, or models of social systems. But the concerns are different from those of the traditional picture. Chaos would reign not because we are fundamentally selfish and unruly and therefore have to create punitive structures to keep ourselves in check, but because we simply do not have the intellectual bandwidth or the psychological acumen to predict, perhaps even make sense of, our fellow human being in the absence of such structures. That is, on this view we can be as motivated to cooperate as we like, but nonetheless be unable to do so because we cannot grasp what other people are doing, intending to do, thinking about what we are doing, and so on. Quite apart from Morton's considerations about entangled social cooperation choices, it is a fact that actions do not come with labels and people don't carry their beliefs, desires, or intentions on their sleeves. In the absence of external structures—such as norms or practices—the interpretive field is wide open. In order to make headway, we have to limit the space of interpretations. Norms do this for us, if nothing else.

What is especially interesting about such views is that they remain agnostic about the moral or political quality of these guiding norms; some more than others. In my view, McGeer is far too optimistic about the rationality undergirding norm-like behavior. Why would we be rushing headlong into a global catastrophe, with our eyes wide open, if we were that reasonable, for instance? Norms need not be particularly morally good, fair, or rational. They just need to enable cooperation in a relatively stable way. Once they are in place, they help narrow the scope of interpretation of others while simultaneously limiting the scope of options for ourselves, since it will be in our interests to follow them even if they are disadvantageous to us personally.

This picture appears to solve the cooperation problem, then, but not in the way we were led to suspect. It is not that you cooperate only because you are able to predict that your partner will also cooperate. This would bring us back to the entangled decision we were faced with before, when dealing with a one-off interaction (common in large-scale societies). In repeated encounters, other tactics can be used to solve the problem, such as adopting a tit-for-tat approach, but here again prediction is not necessary; indeed, it is mostly superfluous. Although the existence of a stable norm increases the predictability of social players, adherence to the norm itself is really what does the job, not this increased predictability (social norms tend not to be conditional on the other person's adherence). In other words, what motivates us to cooperate is the same thing that gives rise to the expectation that the other will cooperate too. But our decision to engage in the particular cooperative act need not be based on this expectation. Indeed, if it were, we wouldn't really be following a norm.

Morton's insight is that social cooperation cannot be based on prediction because of the interdependence of each person's decision. Instead, a social actor needs to adopt a certain approach, then act on it. It may be important, in such cases, to signal clearly to the other person what one's approach is. When a norm is available, all one needs to do is to conform to it irrespective of what the other person does. Of course, should the other person turn out not to cooperate, this might change future choices, but in principle the existence of a norm obviates the need for any kind of prediction about what the other person will do. A different way of putting the same point is to say that a decision to act in accordance with the norm is not itself based on a prediction that the other will act thus-and-so. Instead, it is based on the strategy outlined by the norm.

Calhoun's claim, then, that *predictive* expectations underlie much actual cooperation is problematic. As we have seen, relying on predictions of someone's intentions or actions is difficult, and if the situation approaches something like a one-shot prisoner's dilemma, even more so. Why? Because your decision will be entangled with that of the other person, and the only way out will be to rely on *something else*. And it is here that norms come in, as we have seen. Calhoun, I imagine, might be happy with all this, although her account of a responsible person—beyond praise and blame—relies on these troublesome predictive expectations. Now, it is not unlikely that Calhoun's focus on prediction comes from trying to accommodate the responsibility status of foreigners and, possibly, reprobates. But, of course, in her attempts to accommodate this, she runs into another problem, namely that predicting what others believe, want, intend, or will do is a tremendously complex task, unlikely to underlie the sort of norm-guided action Calhoun is after.

It is certainly true that if social norms are in place, and a society is not itself dysfunctional, then one can predict that people will adhere to them. But Morton's point is that it is not the prediction itself that underwrites the adherence. Instead, I adhere to the norm because the norm constitutes a strategy for solving a particular cooperation problem or, more simply, because it is the norm. Thinking of matters more in terms of proximal psychological mechanisms, I adhere to the norm because I've been socialized to do so, which is of course also what Calhoun suggests. But prediction need not figure in the causal sequence. In fact, it would be a rarity for it to do so, in a well-functioning society, and would typically not be helpful. It seems, therefore, that we should adopt a simpler account according to which the existence of norms themselves, perhaps understood as solution-based approaches to social coordination, do the work.

5.3 Trustworthiness

What I have just argued is open to a simple objection, namely that it just seems to be empirically false. We spend an awful amount of time evaluating people, and the higher the stakes are, the more we do so. For instance, I recently bought a print from a gallery in SoHo. Although by art standards it wasn't that expensive, it was still a significant amount of money. Since I wasn't able to take it with me there and then, I was going to pay for the print and the shipping first and then rely on them sending me the print. I did mentally evaluate the gallery space and the owner before handing over my credit card. Even so, it wasn't the owner I trusted so much as the very structure within which our interaction took place. I considered that it was a gallery I had seen before, that reputation is important for business owners, and that there's always small claims court if they fail to deliver, etc., etc. But there are many other examples. Some years ago, my friend and I went east from Parque Nacional Natural Tayrona along the Colombian coast to a small eco-hotel run by a former professional German cyclist. When the bus dropped us off, two guys on motorcycles offered us a ride to the place. For a fee, of course, which was duly set out on little laminated cards. There were no taxis. It was getting late and we didn't know the way. We did some quick hard thinking and got on the motor cycles. Here, however, the stakes were definitely higher and which social protective structures were in place was simply unclear. Ask any person, and they will tell you of hundreds of encounters in which they considered a potential interaction partner's trustworthiness. So, what's going on?

Just because people unreflectively adopt social roles doesn't mean that they are not also *people*, and we switch back and forth between regarding them in terms of their social roles and in terms of their particular character and situatedness. I can have a pleasant conversation with my barista in Donostia, after which I order a coffee and expect him, not as the person that he is, but as a barista, to make my coffee and wait to get paid until I'm done. During our more personal interactions, folk psychology, as we usually think of it, comes online. But once we enter the personal, concerns about a person's living up to their role or commitments rearise. After all, this particular person might or might not adhere to the norms. Calhoun addresses this issue when she makes the distinction between default and non-default trust, following Annette Baier's influential work on trust (1986). Default trust is what characterizes our attitude to routine social interactions, the context of which I discussed at some length in Section 2. And you don't have to take my—or Baier's or Calhoun's—word for it; there is empirical evidence that supports the idea that we default to trust in social interactions that have the same structure as routine social interactions (Katzir and Posten 2023). Non-default trust, however, is more personal. It might be thought

default trust simply applies to the social, and non-default trust to the private sphere, but this would be a mistake. Many of our expectations derived from the social fabric apply to people in our private sphere, too; our siblings, partners, or children.

Anyway, what happens in situations of non-default trust? We have seen that the standard story, where we base our interactions with others on predicting what they will do, won't work because of the entangled nature of these situations. However, notice that one of the important stipulations of the prisoner's dilemma is that the two prisoners are not able to communicate after they have been informed of the payoffs. This prevents, for instance, their promising each other to cooperate. Robert Frank found that promising is one of the most common strategies people adopt in one-shot prisoner's dilemmas if they are allowed to interact (Frank 1988). And doing so increases the probability that people will, in fact, cooperate. We sometimes make people promise things, but mostly we take it for granted that the other person would not enter into the particular interaction unless they also intended to keep up their end of the bargain. When the stakes are high, of course, we do not. Sometimes we have ceremonies, in which the parties officially swear fealty to one another, such as in a wedding ceremony. Here, the swearing might have as much to do with cooperating as does the swearing of the other person. But leaving those special occasions aside, what about all the cases where we consciously consider whether the other person will keep up their end of the bargain?

This is where trust comes in. And when we think of human cooperation, trust is, of course, where it's at. Trust is the agent-centered complement of expectation (or predictive expectation in Calhoun's case). And you might think that this is the solution to the cooperation problem. We trust the other person, which among other things means we trust them to cooperate in this particular case. But, of course, if the way things work is that we only trust the other person if we know they can be trusted, then we get stuck again. The other person will only trust us if they know we can be trusted, but we can't be trusted unless we know they can be trusted. In the end, someone has to make the first move. Luckily, the evidence suggests that people default to trust, i.e., they are more prone to trust, than not to trust, in a variety of social contexts (Dunning et al. 2014; Katzir and Posten 2023). And, in an interesting twist, people who distrust others are less likely to be trusted by others and are, as a matter of fact, less trustworthy themselves (Weiss, Burgmer, and Mussweiler 2018). Moreover, people who are seen to calculate the pros and cons of cooperating before they commit are trusted less (Jordan et al. 2016).

But, again, this nice quick picture—thanks to a quick perusal of the social psychological literature—belies the fact that we spend a fair amount of time weighing people up. In fact, according to Alan Sillars (2011), who has done extensive work on empathic accuracy, a full third of our thoughts about others concern the nature of our relationships to them. To be fair, Alan Sillars works with people in couples or family therapy, where people are highly likely to monitor each other. Nonetheless, there is plenty of other evidence of social monitoring. For instance, in our impression formation of others, we focus first and foremost on their moral and social traits (Wojciszke 1994; Fiske, Cuddy, and Glick 2007). This suggests that we spend a fair amount of time making psychological ascriptions to others—or figuring out which ones to make—in the context of our social interactions with them. But we do not have to suppose that such ascriptions are for predicting what people will do so that we can decide whether to cooperate with them or not; we use them to determine their attitudes towards us. In philosophy of mind, we talk about character and character traits, but if we turn to work in ethics, for instance, people talk about “good will,” particularly Strawsonians. I am not a Strawsonian myself, but I think there is an intriguing meeting point here between what is regularly being discussed in the responsibility literature and preoccupations we find when we examine people’s folk psychological attributions. In short, part of the function of folk psychology might, in fact, be to determine people’s good will, in general, and their attitudes towards us in particular.

In early criticisms of the classical view of folk psychology, it was noted that belief–desire psychology only constitutes a small part of our folk psychological ascriptions (Andrews 2012). Consequently, traditional conceptualizations of this important ability ignore some of the most important assessments we make of people, namely in terms of their personality traits, such as whether they are kind, friendly, arrogant, smart, beautiful, confident, and so on. Some of the social psychological literature suggests that we make these assessments very quickly upon meeting other people (Wojciszke 1994; Wojciszke, Bazinska, and Jaworsky 1998; Fiske, Cuddy, and Glick 2007). Quite apart from these attributions, we track closely how others relate to us. This need not express itself in conscious ascriptions of propositional attitude states, however, but in how we feel about others. Do they make us feel uncomfortable, angry, ashamed, happy, smart, interesting, or rejected? Notice that this assessment is the result of two things; the emotion itself and its supposed cause (namely the other person).

Moving on to the issue of most concern when reflecting on the topics discussed by Calhoun in this lecture, there are a variety of things that people seem sensitive to when they assess the trustworthiness of potential cooperators in non-default

trust contexts. I find particularly intriguing the evidence coming out of Molly Crockett’s lab (Everett, Pizarro, and Crockett 2016; Everett et al. 2018). Her and her collaborators report that people put less trust in people who make utilitarian choices in the type of tough moral dilemmas that proliferate in the psychological literature on morality. This is true even of people who make utilitarian choices themselves (Everett et al. 2018). Such dilemmas include, for instance, the footbridge version of the trolley dilemma (also known as a “sacrificial dilemma”). Here, we are faced with the choice of either letting a runaway trolley continue onto a track with five hikers, who will be killed as a result, or throwing the fat man next to us on the footbridge onto the tracks, thereby stopping the train and saving the five hikers. However, as some of their studies show (Everett et al. 2018), it may not be the impartial benevolence of utilitarianism that people have a beef with, but with the instrumentalist attitude shown towards the welfare of the sacrificial victim. In scenarios where the sacrificial victim is purely the means to the end of saving the five hikers, we find the pattern discussed above. However, when the sacrifice of one person is a mere side-effect of saving the five hikers, people’s utilitarian choices do not prejudice others against cooperating with them. In other words, treating people merely as means to an end is generally frowned upon, even when doing so saves more people. This makes sense from a cooperation standpoint. Someone who is willing to treat another person as means to an end of optimizing some benefit might equally do so with us in a collaborative ambit.

The bias against act utilitarians—in the limited situations where this choice tendency has been explored—indicates that we are, indeed, sensitive to the information we have about other people when we are faced with a choice to cooperate with them, at least in situations not scaffolded by social or moral norms and structures (i.e., non-default trust situations). This is true even if our default is to trust others. However, together with the fact that individual prediction of cooperation cannot play the role in human cooperation we have been led to believe, the extant evidence from social psychology supports the idea that we do make evaluations of trustworthiness, and that they sometimes form the basis of our decisions to cooperate, at least in certain experimental situations such as the trust game (where you have the choice of transferring money to your partner, which will then be doubled, and where in order not to be worse off you have to trust your partner to transfer back to you at least as much as you transferred to him in the first place).

5.4 Conclusion

Compliance responsibility, argues Calhoun, enables social life because it gives people a reason to suppose that others will, as a matter of fact, adhere to the common

norms. Without such reasons, there would not be sufficient reason for *them* to adhere to such norms either. Or so the thought usually goes. After all, my cooperating with someone who does not cooperate with me tends to lead to a net loss to me and a net gain to them. Therefore, I am better off not cooperating with them if I have reason to think they will not cooperate. This “cooperation problem” has been modeled in various ways in decision theory, most famously by the prisoner’s dilemma. During the eighties and nineties in philosophy of mind, it used to be thought that the solution to the cooperation problem was folk psychology or theory of mind. This would allow us to predict other people’s intentions, and thereby their actions, which would make it possible for *us* to plan our own actions. Calhoun, too, seems to think that predictability is important, hence her “predictive expectations.” But, as we have seen, relying on prediction to solve the prisoner’s dilemma gets us nowhere. Indeed, it is obvious that if I need to have reasons to suppose that others will adhere to the common norms before I adhere to them, then they, too, need to have reasons to suppose that *I* will adhere to common norms before they adhere to them, thus getting us into the regress problem we discussed earlier. We cannot make our decisions premised on predictions about what others will do. Instead, we must choose a good strategy and act on it. One might think norms of behavior as developments of such strategies.

Following a thread coming from theorizing about folk psychology, or theory of mind, then, provides support for Calhoun’s general idea of norms adherence being central to social life. But it also suggests that she should not rely on *predictions* of other people’s norm adherence. It might then be thought that perhaps a predictive route goes through something other than intentions, such as someone’s character. I think it is quite likely that we use such an approach when it comes to choosing friends, partners (romantic and business), and so on. Of course, when it comes to the sorts of routine interactions we have with people in our everyday lives, they are unlikely to play much of a role in our actual decisions to adhere to social norms. Given Calhoun’s emphasis on trust, I wonder how much she actually wants to rely on “predictive expectations.” Whatever the answer turns out to be, I have given reasons to think that prediction cannot underlie social cooperation on a relatively common way of thinking about such cooperation.

It would be absurd, of course, to maintain that folk psychological predictions *never* play any role in interactions with others. My point has been that predictions won’t solve the cooperation problem, as it has typically been modeled. There are reasons to think that predictions play much less of a role than explanations in psychological ascriptions, and that it is the influence of the idea that folk psychology is like a

scientific theory that has given that impression (cf., e.g., Andrews 2012). Having said that, however, I don't see any use in denying that predictions of others' intentions or actions sometimes play a role in social life.

In the context of thinking about the role of psychological ascriptions, I also referenced work suggesting that our very ascriptions are norm-guided and norm-guiding. Why? Because the field of possible explanations of others is too large otherwise. This is another way to look at norms. It's not just that they are there to regulate our social interactions, enabling the kind of cooperation that has given rise to large-scale societies, but also that they are needed for us to become understandable to each other. This is another way of stressing the idea that folk psychological predictions cannot themselves solve the cooperation problem.

The other issue I raised has to do with norm optimism. Moral and social norms in the wild, as opposed to in the Platonic structures philosophers often discuss, are pretty ambiguous. I here understand moral and social norms as those that hold, as a matter of empirical fact. One might regard such moral norms as morally reprehensible, in the absolute sense (although who gets to decide, or "discover" absolute right or wrong is a question fraught with difficulties). In my view, the two senses of "moral" are not distinguished often enough. This may be why so many philosophers view norms in a rather starry-eyed fashion. If we can somehow manage to institute "true" moral norms (in the absolute sense)—and so bring the empirical norms into line with the absolute ones—one might find it hard to object to norm optimism (although I rather suspect that such optimism relies on the doubtful assumption that one's pet moral beliefs will invariably turn out to reflect absolute ones). Leaving that sticky issue aside, it is an empirical fact that norms on the ground, as it were, are often not very beneficial to a great number of individuals to whom they apply. I don't think this point can be stressed enough. Norms force us to do what we are not otherwise inclined to do; suffer insults, rape, misery, and starvation. Norms deprive us of our freedom to be who we are, to do what we want, to express our opinions, to live with dignity, to experiment with ways of living, and even to wear what we want. Every day, all over the world, people are shunned, harmed, or killed as a result of prevailing socio-moral norms. Sometimes norms prevent people from harming or killing others, but often they facilitate doing just that. Norms exculpate harm, quite possibly as often as they prevent them. It's just that the person who is harmed must be an enemy (as the recent attack by Hamas on Israelis amply demonstrates), must have different religious beliefs (cf. wars of religion), deserve it (cruel punishment), must not really be a person (as in much genocidal propaganda), and so on. Norms have many sides, and we should not forget the ugly ones.


Civilization comes at a cost, of course (Freud [1930] 1961). It reduces an individual's freedom and the opportunities she is free to pursue. The responsible man or woman—or however the person chooses to classify themselves—leaves a world of opportunity behind and takes up the yoke of their particular culture. This may, or may not, be to his or her individual advantage. Once, however, someone is born into a certain system, the balance of power almost invariably makes it the case that it is to their advantage to comply with the norms. And so, given the tawdry history of moral and political norms, we ought to approach the issue of normativity in human life with more caution. It might, for instance, turn out that power and power imbalance is a design feature of norms in the wild.

Anyway, it is time to release you from my preoccupations and let you take in another perspective on Calhoun's fascinating lectures. Thanks for reading!

5.5 References

- Andrews, Kristin. 2012. *Do Apes Read Minds? Toward a New Folk Psychology*. Cambridge, MA: MIT Press.
- Anthony, Andrew. 2021. “‘We Showed It Was Possible to Create a Movement from Almost Nothing’: Occupy Wall Street 10 Years On.” *Guardian*, 12 September, 2021.
- Axelrod, Robert. 1984. *The Evolution of Cooperation*. New York: Basic Books.
- Baier, Annette. 1986. “Trust and Anti-Trust.” *Ethics* 96: 231–60.
- Bellware, Kim. 2022. “Georgia Supreme Court Reinstates Six-Week Abortion Ban.” *Washington Post*, November 23, 2022.
- Binmore, Ken. 1994. *Game Theory and the Social Contract. Vol. 1: Playing Fair*. Cambridge, MA: MIT Press.
- Calhoun, Cheshire. 2004. “An Apology for Moral Shame.” *The Journal of Political Philosophy* 12: 127–46.
- Ciabattari, Jane. 2023. “Why Is the US Banning Children’s Books?” *BBC News*, 26 May, 2023.
- de Beauvoir, Simone. (1949) 2011. *The Second Sex*. Translated by Constance Borde and Sheila Malovany-Chavallier. New York: Vintage.
- de Waal, Frans. 1982. *Chimpanzee Politics*. Baltimore, MD: Johns Hopkins University Press.
- Dunning, David, Joanna E. Anderson, Thomas Schlösser, Daniel Ehlebracht, and Detlef Fetchenhauer. 2014. “Trust at Zero Acquaintance: More a Matter of Respect than Expectation of Reward.” *Journal of Personality and Social Psychology* 107: 122–41.
- Everett, Jim A. C., David A. Pizarro, and M. J. Crockett. 2016. “Inference of Trustworthiness from Intuitive Moral Judgments.” *Journal of Experimental Psychology: General* 145: 772–87.
- Everett, Jim. A. C., Nadira S. Faber, Julian Savulescu, and Molly J. Crockett. 2018. “The Costs of Being Consequentialist: Social Inference from Instrumental Harm and Impartial Beneficence.” *Journal of Experimental Social Psychology* 79: 200–16.
- Fiske, Susan T., Amy J. C. Cuddy, and Peter Glick. 2007. “Universal Dimensions of Social Cognition: Warmth and Competence.” *Trends in Cognitive Sciences* 11: 77–83.
- Fodor, Jerry. A. 1989. *Psychosemantics: The Problem of Meaning in the Philosophy of Mind*. Cambridge, MA: MIT Press.
- Frank, Robert H. 1988. *Passions Within Reason*. New York: W. W. Norton & Co.
- Freud, Sigmund. (1930) 1961. *Civilization and Its Discontents*. Translated by James Strachey. New York: Norton & Norton.
- Graeber, David and David Wengrow. 2021. *The Dawn of Everything: A New History of Humanity*. London: Allen Lane.
- Godfrey-Smith, Peter. 2004. “On Folk Psychology and Mental Representation.” In *Representation in Mind: New Approaches to Mental Representation*, edited by Hugh Clapin, Philip Staines, and Peter Slezak, 147–62. Amsterdam: Elsevier Publishers.
- Hobbes, Thomas. (1668) 1994. *Leviathan*. Edited by Edwin Curley. Indianapolis, IN: Hackett Publishing.
- Hrdy, Sarah Blaffer. 2009. *Mothers and Others: The Evolutionary Origins of Mutual Understanding*. Cambridge, MA: MIT Press.
- Jordan, Jillian J., Moshe Hoffman, Martin A. Nowak, and David G. Rand. 2016. “Uncalculating Cooperation Is Used to Signal Trustworthiness.” *PNAS* 113: 8658–63.
- Katzir, Maayan and Ann-Christin Posten. 2023. “Are There Dominant Response Tendencies for Social Reactions? Trust Trumps Mistrust—Evidence from a Dominant Behavior Measure (DBM).” *Journal of Personality and Social Psychology*.

- Kohlberg, Lawrence. 1984. *Essays on Moral Development*. New York: Harper & Row.
- Maibom, Heidi. 2003. "The Mindreader and the Scientist." *Mind & Language* 20: 237–57.
- Maibom, Heidi L. 2007. "Social Systems." *Philosophical Psychology* 20: 557–78.
- . 2010. "The Descent of Shame." *Philosophy and Phenomenological Research* LXXX: 566–94.
- McGeer, Victoria. 2007. "The Regulative Dimension of Folk Psychology." In *Folk Psychology Re-Assessed*, edited by Daniel D. Hutto and Matthew Ratcliffe, 137–56. Dordrecht: Springer.
- Morton, Adam. 1996. "Folk Psychology Is Not a Predictive Device." *Mind* 105: 119–37.
- . 2003. *The Importance of Being Understood: Folk Psychology as Ethics*. London: Routledge.
- Nietzsche, Friedrich. (1887) 1989. *On the Genealogy of Morals*. Edited and translated by Walter Kaufmann and Reginald John Hollingdale. New York: Vintage.
- NPR Staff and Wires. 2011. "Occupy Wall Street Inspires Worldwide Protests." *NPR News*, October 15, 2011.
- Raihani, N. J. and R. Bshary. 2011. "Resolving the Iterated Prisoner's Dilemma: Theory and Reality." *Journal of Evolutionary Biology* 24: 1628–39.
- Richerson, Peter J. and Robert Boyd. 2005. *Not by Genes Alone: How Cultures Transformed Human Evolution*. Chicago: University of Chicago Press.
- Sautter, John A., Todd A. Brown, Levente Littvay, Alberta C. Sautter, and Brennen Bearnes. 2008. "Attitude and Divergence in Business Students: An Examination of Personality Differences in Business and Non-Business Students." *Electronic Journal of Business Ethics and Organization Studies* 13: 70–78.
- Sillars, Alan. 2011. "Motivated Misunderstanding in Family Conflict Discussions." In *Knowing When—and When Not—to Understand Others*, edited by Jessi L. Smith, William Ickes, Judith A. Hall, and Sara D. Hodges, 193–212. Hauppauge, NY: Nova Science.
- Spaulding, Shannon. 2018. *How We Understand Others: Philosophy and Social Cognition*. New York: Routledge.
- Staver, Anna. 2021. "House Republicans Introduce Bill to Ban Teaching of Critical Race Theory in Ohio." *The Columbus Dispatch*, 25 May, 2021.
- Sterelny, Kim. 2012. *The Evolved Apprentice: How Evolution Made Humans Unique*. Cambridge, MA: MIT Press.
- Watkins, Morgan. 2023. "After Yearlong Fight, a Near-Total Abortion Ban is Going into Effect in Indiana." *NPR News*, August 1, 2023.
- Weaver, Matthew. 2023. "Kathleen Stock Says She Is a 'Moderate' as Protests Planned over Oxford Debate." *Guardian*, May 29, 2023.
- Weiss, Alexa, Pascal Burgmer, and Thomas Mussweiler. 2018. "Two-Faced Morality: Distrust Promotes Divergent Moral Standards for the Self versus Others." *Personality and Social Psychology Bulletin* 44: 1712–24.
- Wojciszke, Bogdan. 1994. "Multiple Meanings of Behavior: Construing Action in Terms of Competence or Morality." *Journal of Personality and Social Psychology* 24: 1251–63.
- Wojciszke, Bogdan, Roza Bazinska, and Marcin Jaworsky. 1998. "On the Dominance of Moral Categories in Impression Formation." *Personality and Social Psychology Bulletin* 24: 1251–63.
- Zawidzki, Tadeusz W. 2008. "The Function of Folk Psychology: Mind Reading or Mind Shaping?" *Philosophical Explorations* 11: 193–210.
- Yurcaba, Jo. 2023. "DeSantis Signs 'Don't Say Gay' Expansion and Gender-Affirming Care Ban." *NBC News*, May 17, 2023.



6

RESPONSIBILITY: EXPECTED, TAKEN, RECOGNIZED

Gunnar Björnsson

Stockholm University

6.1 Calhoun on Responsible Persons

Plausibly, to be responsible is to satisfy all general preconditions for being fittingly held responsible. Philosophers of responsibility have thus paid avid attention to the nature and preconditions of the most conspicuous practices of holding one another responsible: of blame for violations of moral expectations.

In her Descartes Lectures, Cheshire Calhoun asks us to widen our gaze.⁶² To understand responsibility, we should also take into account our basic trust that others will fulfil at least the most basic normative expectations, and our predictive expectation that people will do good things that they are not required to do. Likewise, we should take into account our positive reactions of appreciation and gratitude towards those who do, reactions that Calhoun takes to be importantly different from negative reactive attitudes. And we should take into account our negative reactions to *not* being held accountable, expected to do the right thing, or asked to contribute what isn't normatively expected of us.

Together, Calhoun argues, these practices, reactions, and attitudes reveal that we operate with a default assumption that people are *responsible persons*. A responsible

⁶² Others have similarly argued for a more diverse diet of examples. See in particular Shoemaker forthcoming.

person, on her proposal, is someone who is (i) *accountable*—able to live up to normative expectations⁶³—(ii) *compliance responsible*—in fact living up to basic normative expectations, making them also predictive expectations⁶⁴—and (iii) a *responsibility taker*—at least sometimes disposed to take the initiative to do good things that are not normatively expectable. Furthermore, being a responsible person is understood as valuable, as is being recognized as one.

Considering a wider range of related phenomena is often revealing, not only because the new phenomena are themselves interesting, but also because they make us see old phenomena with fresh eyes. In this case, the widened gaze helps us to think of responsibility and our responsibility practices in finer detail. It is in that spirit that I consider Calhoun’s positive account, asking further questions about the notion of a responsible person. Specifically, I will ask whether she is right that we operate with the default assumption that people are responsibility takers in her sense. Contrary to Calhoun’s proposal, I will suggest that, once we take into account the full range of standard normative expectations, it is doubtful that our predictive default expectations of what we naturally describe as responsibility taking outstrips what is normatively expected. In addition, I propose a way of nevertheless making good on Calhoun’s suggestion that accountability, compliance responsibility, and contributions to the common good that merit gratitude are all aspects of responsibility. Finally, I suggest that what positive reactive attitudes reveal about their targets is not that they are responsibility takers, but that they are weight-givers subject to “balancing norms,” which call on us to care about giving people and certain other values a certain comparative weight over time.

6.2 Do We Assume, by Default, that People Are Responsibility Takers?

Calhoun is clearly right that “our social practices are pervasively structured on the presumption that social participants have the capacity and disposition to elect to promote the good that underwrites those practices in non-required ways” (68–69). As she points out, a great many organizations, including charitable organizations,

⁶³. The ability need not be the ability to directly grasp or live up to normative expectations: sometimes we need assistance from the surrounding community, in line with Victoria McGeer’s (2019) “scaffolding account” of the reactive attitudes and corresponding account of responsibility.

⁶⁴. *Basic* normative expectations include constitutive norms of practices and norms that are socially understood as matters of common decency rather than more elevated normative expectations (32). Exactly how to draw the distinction between the latter two categories is left open (34, n. 30), but one possible test, inspired by sociologist Harold Garfinkel’s (1964) “breaching experiments,” is in terms of “how difficult it would be for an average social participant to bring themselves to violate the rule. ‘Don’t negatively comment on other’s purchases’ might be among those. Other examples of norms of common decency would include not only prohibitions on intentionally and severely harming others, but also the expectation that a ‘hello’ greeting will be returned rather than ignored and, in the U.S., the expectation that guests will not smoke in one’s house unless given permission.” (34)

churches, clubs, professional organizations, and political advocacy groups rely, to various extents, on people voluntarily contributing money or labor, and workplaces frequently rely on people voluntarily taking on tasks. In our own profession, the voluntary refereeing of journal submissions, project funding, tenure, and promotions plays a crucial role. Calhoun is also clearly right that such voluntary contributions call for—and standardly receive—gratitude and appreciation. Based on these observations, she concludes that

the default presumption is that social participants are sufficiently capable of appreciating the goods served by norm-structured practices, have sufficient motivation to adopt some of those goods as personal ends, and have sufficient cognitive capacity to see how those goods might be promoted in non-required ways for it to make sense to organize social life so that only some promotion of the good is normatively expected and much left normatively optional. (69)

Differently put, she concludes that our notion of a responsible person includes that of being a *responsibility taker*, where taking responsibility is electing “to do things that are not normatively expectable” (59).⁶⁵

This conclusion, I think, should be rejected. I will suggest that, to the extent that we presume by default that social participants are motivated to adopt and promote some social ends, this is already implicit in the assumption that people are accountable, or able to live up to normative expectations, and the assumption that they actually live up to normative expectations.⁶⁶ The suggestion builds on the idea that normative expectations go beyond expectations of certain kinds of fairly well-defined actions, such as expectations to return greetings, not to take what is not one’s own, or to keep promises. Crucially, they also include expectations to care about certain ends, expectations revealed by negative reactive attitudes when agents fail:

NORMATIVE EXPECTATION: People are normatively expected to be willing to contribute to the common good and to help others to some extent, giving the ends of others and shared practices some weight.

⁶⁵ She also notes that we would often find it insulting not to be included in requests for help, when this would suggest that we are unable or unwilling to contribute: it is understood as valuable to be seen as willing and able to contribute.

⁶⁶ The latter condition goes beyond the assumption of what Calhoun calls *compliance responsibility*, which only involves living up to *basic* normative expectations. Exactly how to delimit basic expectations is a little unclear (see footnote 64), leaving me uncertain about whether the normative expectations that I appeal to are basic or not. However, this does not matter for my purposes, as I am concerned with whether people are expected to contribute beyond what is generally normatively expectable, not just beyond basic expectations.

In particular, people are normatively expected to help when asked to help, unless the requests are illegitimate or they have sufficient reason not to help.

Based on NORMATIVE EXPECTATION, I will further claim that:

PREDICTIVE EXPECTATION: Individuals are not in general predictively expected, by default, to do more for the common good or others than what is normatively expected.

Though PREDICTIVE EXPECTATION is an empirical claim, I will just rely on the reader's assessment of its plausibility, based on personal experience. My focus will be on supporting and explaining normative expectation.

To understand my target here, it is important to distinguish two claims:

NON-DISTRIBUTIVE PREDICTION: We predictively expect, by default, that in a sufficiently large social group, some individuals will do more for the common good, or for others, than what is normatively expected.

DISTRIBUTIVE PREDICTION: We predictively expect, by default, of each individual that they will do more for the common good or others than what is normatively expected.

Given natural variation in people's dispositions to contribute beyond what is normatively expected, NON-DISTRIBUTIVE PREDICTION is no doubt correct for large enough groups of participants. (For the same reason, we might presume, by default, that in a sufficiently large group there will be occasional failures to live up to normative expectations). What Calhoun alleges, and what I remain unconvinced about, is DISTRIBUTIVE PREDICTION. It is not clear to me that we predictively expect of

individuals, by default, that they will contribute beyond what we normatively expect them to do.⁶⁷

It is also important to distinguish normative expectation from the claim that we are normatively expected to contribute to the common good or help others *in specific ways*. Calhoun's examples make the solid point that many of the specific ways in which people contribute to the common good and help others are not normatively expected: I'm not normatively expected to take on *this* refereeing task, or contribute to *this* charity. But such specific contributions are also not predictively expected. Moreover, the examples are compatible with the existence of a general normative expectation that people contribute *in some way or other* to these goods, beyond the specific ways that are normatively expected.⁶⁸ Part of what I suggest is that we are under such general normative expectations to contribute. If we are, then our normative expectations might line up with our predictive expectations.

The notion of a normative expectation is obviously critical for the interpretation and plausibility of NORMATIVE EXPECTATION. Calhoun ties the notion to that of an obligation, or to what one may "properly demand of others within a shared normative community" (53). But what is it to *demand* something of others, in the relevant sense? In light of the role that normative expectations play in Calhoun's understanding of accountability, I take relevant demands to include those implicit in the negative reactive attitudes of indignation and resentment. Exactly how to understand such demands might not matter, but I will assume that they involve a disposition to treat their targets less favorably until they have taken suitable and

^{67.} Similarly, what Calhoun alleges is not merely that, as a matter of statistics, we can, by default, expect everyone to at some point contribute beyond what is normatively expected of them, just as we can expect everyone, by default, to at some point fall short of expectations. The responsibility-taking that she thinks that we expect of others, by default, is supposed to be, in aggregate, a very significant part of social life, not mere random deviations from the norm. A complication in understanding Calhoun's argument concerns the notion of a *default assumption*, an assumption that holds pending evidence to the contrary. Obviously, our normative as well as predictive expectations on children and adults differ, and various stereotypes might lead people to expect more from some than from others: more empathetic attention from women than from men say, and more pro-sociality from white compatriots than from immigrants of color. On the face of it, these are differences in default assumptions. Alternatively, signs that someone belongs to a certain category (child, adult, child, woman, man, white, person of color) might perhaps be seen as modifiers of a basic set of universal default assumptions. I set these issues aside, as my main arguments are independent of how they are resolved.

^{68.} Many normative expectations leave open how they are satisfied: obviously, expectations not to lie, steal, or kill standardly leave open a great variety of ways in which one may avoid lying, stealing, or killing, while expectations to do specific positive actions leave open how exactly these actions are performed. I take Calhoun's repeated claim that we predictively expect people to contribute in ways that are normatively optional to go beyond this triviality: what she suggests is that we expect people to contribute *to an extent* that isn't normatively expected.

sufficient steps to amend the situation, by apology and efforts to repair or compensate for the harm their culpable wrongdoing might have inflicted. On this understanding of *NORMATIVE EXPECTATIONS*, normative expectation says that if people are not to some extent willing to help others, contribute to the common good, and respond positively to requests for such actions, they would be subject to indignation and resentment, which involves some tendency to treat them less favorably.

6.3 Requiring Reasons, Evaluative Autonomy, Balancing Norms

Understood along the lines just suggested, normative expectation strikes me as plausible. Calhoun is right that it is often consistent with *NORMATIVE EXPECTATIONS* not to act on requests for aid or to promote social ends, instead pursuing one's private projects. But I take it that we do react with (perhaps mild) indignation to someone who never contributes in spite of ample opportunity to do so at low cost to themselves, and that we are prone to hold this against them, taking ourselves to be (*pro tanto*) justified in showing them somewhat less goodwill than we would to those contributing to a normal extent.

In addition, I suggest, normative expectations on our readiness to help others are extensive. Suppose that someone carrying grocery bags is struggling to open a door and risks dropping the bags. Suppose further that I don't help in spite of noticing their predicament, not because I'm occupied or think that I have something better to do, but because I just don't care. Then guilt on my part would be fitting. And should you learn about what I did, and why, you could fittingly be at least mildly indignant with me, and retract some amount of goodwill until I at least expressed some regret or showed, in other ways, that I care enough about the plight of strangers enough to give them a hand when good opportunities arise.⁶⁹ Or suppose that a stranger asks you for directions but you shake your head and move on, not because you thought you had something more important to do, but because you just attributed no significance to their request. Then negative reactive attitudes would again be fitting. Perhaps these attitudes should be relatively mild, as you weren't under a particularly stringent duty to help; if you had been somewhat short of time, or about to make a phone call, or even just deep in thought, turning down the request would have been fine. But the point is that another person's reasons, in particular reasons they invoke in asking you for help, often seem to be *your* reasons for action, to be set aside only if there are sufficiently strong countervailing or undermining reasons.⁷⁰

⁶⁹. Would I also be violating a *basic* normative expectation? On the proposal reproduced in footnote 64, this strikes me as plausible.

⁷⁰. In the words of Christine Korsgaard (1996, 140): "If I call out your name, I make you stop in your tracks. ... Now you cannot proceed as you did before. Oh, you can proceed, all right, but not just as you did before. ... By calling out your name, I have obligated you. I have given you a reason to stop."

Such requiring reasons, and resulting (perhaps weak) *pro tanto* obligations, straightforwardly account for a phenomenon that Calhoun discusses but struggles somewhat to explain: when we decline requests, we feel the need to offer what looks like excuses or justifications.⁷¹

Importantly, the fact that the interests and requests of others provide us with *pro tanto* obligations is compatible with extensive moral freedom. Not all interests and requests have the same, or indeed any, authority: there is less or no force to malevolent interests, or to requests obviously interrupting sufficiently important ongoing activities that require continuous focus. Furthermore, the strength of the reasons provided plausibly depends on whether one is uniquely well placed to promote them or provide the assistance needed, whether one is the source of the problem that needs solving or has benefited from it, whether one has agreed to attend to the sort of problem at hand, and what social ties one has to whoever needs assistance and makes the request.⁷²

To this I want to add something that I suspect is crucial in understanding moral freedom: to a significant extent, the relative normative importance of an agent's various interests and concerns is *up to them*. At a first pass, something like the following seems plausible:

EVALUATIVE AUTONOMY: If you think that an activity which competes with the common good or the ends of others is quite important, then it *is* quite important and can outweigh fairly weighty competing reasons. Conversely, if you think that it doesn't matter much in comparison, then it doesn't.⁷³ This principle presumably requires qualification: if nothing else, assessments of importance might lack legitimacy to the extent that they reflect ill will or moral indifference.⁷⁴ Still, as long as it is basically on the right track, it ensures considerable moral freedom, even given the

⁷¹. See pages 67–68. Calhoun suggests that the offering of excuses and justifications “sends the message: ‘I understand the good at stake, and my declining shouldn’t be taken as evidence that I’m not a responsibility taking kind of being who is unable to appreciate the good and elect to promote it.’ ... It’s not one’s virtue that needs defending, but one’s default status as a responsible person” (68). I agree that it is not virtue that needs defending. But it is also not one’s default status as a responsible person. Rather, it is the validity of the request and the standing of the person making the request as a reason-giver that needs to be acknowledged. Not acknowledging that standing would reveal one as not just lacking in virtue (understood as excellence), but as falling short of normative expectations.

⁷². For an overview of sources of special obligations and responsibilities, see Björnsson and Brülde 2017.

⁷³. Though I cannot discuss this here, one might also think that individuals also have some authority over the comparative importance of shared ends. The idea that personal projects have importance that competes with what is impersonally important is familiar from Williams 1981.

⁷⁴. If our reasons are relative to what we take to be important, cases where our view of what is important changes over time might pose difficult questions. For a helpful recent discussion of corresponding issues for attitude-sensitive accounts of wellbeing, see Bykvist 2022.

normative importance of the common good and the ends of others, and the normative expectations on us to respond to this importance. Yes, it is important to help, but so are the personal projects and activities that we take to be important.

Earlier, I insisted that we are under extensive normative expectations to help. I have now claimed that we can non-culpably avoid helping if we judge that whatever we would have to give up to help is sufficiently important to us. Have I then taken away with one hand what I gave with the other? Not so. We still violate normative expectations when we fail to help when helping wouldn't sacrifice anything of sufficient importance. Moreover, as I have already briefly suggested, there are limits to the extent to which we can prioritize our own activities and projects: *we are normatively expected to provide some help over time, given suitable opportunities.*

Let me now add some structure to the latter suggestion.

The basic idea is this: It is morally important not only to behave or avoid behaving in certain specific ways (to do as one promised, to conform to specific norms of politeness; to not lie, steal, harm, kill). On top of these familiar deontic concerns,

comparative weight: It is morally important to give persons a certain comparative weight over time. Giving weight to someone, in the sense that I'm after here, is *investing resources*—cognition, time, energy, property—in *promoting their interests or acting on their point of view*. The amount of weight given to someone is a matter of both the extent to which their interests are furthered, and their point of view acted on, and the amount of resources actually invested in this.⁷⁵

For illustration, consider a group of friends deciding what to do together. Intuitively, everyone should be given *an equal say*. As not everyone can speak at the same time, one person might justifiably talk over the others if they then proceed to listen: what is important is that everyone has an equal say *over the span of the deliberation*. Likewise, everyone's equally strong preferences should ideally be given equal weight by the group in their decisions. Nevertheless, as preferences diverge the group might justifiably act on the preferences of some of their members. If they do

⁷⁵. COMPARATIVE WEIGHT and related principles introduced below should be seen as simplified models. To mention just one complication: not all weight-giving matters, or matters equally. Resources invested in acting on an adult's interest but against their will might not count; nor might resources invested in acting in normatively expected ways that others rely on as a matter of course: in respecting their property or bodily integrity, or keeping a promise to them, say, when the costs of doing so are clearly not prohibitive. The latter constraint might involve expectations involved in what Calhoun talks about as *default or basic* trust.

so, however, and especially if the same members repeatedly draw the long straw, it becomes increasingly more important that the group also acts on the preferences of other members. Intuitively, the equal importance of the members calls for the group's actions to give them equal weight over some relevant period of time.

In these cases, it is important that the *group* gives *equal* weight to its members over time. It can also be important that *individuals* give equal weight to individuals over time, or *unequal* weight, as the case may be. For example, it might be important that a parent gives roughly equal weight to their two children over time, more weight to their own children than to the neighbor's, and not too much weight to themselves compared to what they give to their children.

I now suggest that the pattern of these examples further extends to the weight given overall to others and to shared ends:

COMPARATIVE WEIGHT (EXTENDED): It is important that we give others and the common good a certain weight over time, compared to the weight we give to ourselves. The importance of balanced weight-giving over time, I further suggest, is reflected in normative expectations:

CARING ABOUT COMPARATIVE WEIGHT: We are normatively expected to care about others and the common good in a way that involves caring about giving them a certain comparative weight over time. To *care about something* in the sense that I have in mind here is to be disposed to notice what promotes or prevents it, and to be motivated to act on such information. If we fail to give others and the common good the right comparative weight in spite of having been given opportunities to do so, this normally means that we don't care about them as is normatively expected of us. This, I suggest, is what underpins indignation and resentment towards those who again and again prioritize their own interests in spite of opportunities to contribute to something of shared importance. It also makes straightforward sense of Calhoun's observation of a kind of reply naturally offered in response to requests for help: "I already gave" (67). Having already contributed enough to ends of the relevant nature, one violates no normative expectation on comparative weight-giving over time in preserving one's resources for other ends.

The resulting picture has three main components. First, the common good and interests and preferences of others make demands on us. Second, these demands leave us with considerable normative autonomy, as we are not normatively expected to act on them if that would require sacrificing other values of sufficient importance,

and as the importance of our personal projects and activities is partly determined by our own judgments. Third, this normative autonomy is restricted by what we might call “balancing norms,” norms requiring that we care about giving a certain comparative weight to others and to the common good over time.

This three-part picture provides additional background to the premises behind the central contention of Section 2, repeated here:

NORMATIVE EXPECTATION: People are normatively expected to be willing to contribute to the common good and help others to some extent, giving the ends of others and shared practices some weight. In particular, people are normatively expected to help when asked to help, unless the requests are illegitimate or they have sufficient reason not to help. PREDICTIVE EXPECTATION: Individuals are not in general predictively expected to do more for the common good or others than what is normatively expected. Again, as far as I can tell, we do not predictively expect, by default, that individuals contribute more than is normatively expected of them, taking into account normative expectations of the sorts that I have sketched here. In Calhoun’s terms, we do not generally predictively expect others to be *responsibility takers* in addition to complying with normative expectations.⁷⁶

One might worry that CARING ABOUT COMPARATIVE WEIGHT and the overall picture that I have painted of our duties towards one another portrays our ethical lives as implausibly calculating. I will address three versions of this worry.

First, one might worry that the picture gives undue weight to balancing ideals and balancing norms, as opposed to the people involved and their interests. What matters when someone needs help, or when a parent has focused their attention on one of their children at the expense of the other, is that the person in need gets help and that the interests and point of view of the other child are given weight, not that some impersonal value of balance is achieved. Worrying about the balance, it might seem, is having one thought too many, or focusing on the wrong thing.

To see why this worry is misplaced, notice that CARING ABOUT COMPARATIVE WEIGHT doesn’t understand caring about the balance of weight-giving as separate from caring about the people involved. It is because a good parent cares about each of their

⁷⁶ The structure also accounts for phenomena that are often explained with reference to “imperfect duties”. Such duties are said to involve latitude—we can decide not to act on them on particular occasions (see e.g. Hill 1971)—or understood as requirements to adopt certain ends rather than performing a specific action (see e.g. Stohr 2011), or as requirements that we do enough over time (see e.g. Pummer 2023, ch. 6).

two children that they care about each being given the right comparative weight over time, not because the parent cares about some independent value of equality. Similarly, it is because a good person cares both about the common good, and about their own projects, that they care about giving them a certain comparative weight over time.

Second, one might worry that, on the proposed picture, we do not have to care about others now if we have already helped or know that we will attend to them later. This, though, is not an implication of the view. If it is important that two values are given a certain comparative weight, it does not follow that if one has been given weight at the expense of the other, it now lacks importance. What follows is that if the two values call for conflicting actions, it becomes comparatively more important to give weight to the latter. “I already gave” can explain why I am justified in prioritizing personal projects rather than giving more, without implying that the cause at hand no longer matters to me.

Third, one might worry that the view implausibly implies that there are precise balancing calculations to be had, somehow backed up by moral reality, as opposed to a messy social context where norms and the weight of needs are under constant negotiation. But such metaphysical assumptions are not part of the picture. What I have said is compatible with balancing norms being indeterminate to various degrees, and it is hard to avoid the conclusion that they are. Though we might have a sense of the amount of weight given to someone, the measurement is necessarily messy. Giving someone weight, I’ve said, is investing resources such as cognition, time, energy, or money into promoting their interests or acting on their point of view. It seems wildly implausible that the use of different resources, promoting very different interests—often subject to evaluative autonomy—or acting in different ways on various aspects of someone’s point of view can be summed up with any precision. Given the messiness of what should be balanced over time, it is also hard to believe that balancing norms will come with much precision. Moreover, for all I have said, the norms in question might be to a significant extent socially constructed in some sense or other. What matters for the picture I have presented is that there are norms with the content needed to explain the phenomena that we are interested in; their metaphysical status is of unclear relevance.

6.4 Responsible Behavior, Responsible Persons, and Default Assumptions

On Calhoun’s picture, predictively expecting people to be *responsible persons* in the relevant sense is taking them to be (i) accountable, (ii) compliance responsible,

and (iii) responsibility takers. In effect, I have questioned whether the last of these components adds to the second.⁷⁷ What we are predictively expecting of people with respect to helping others, responding to requests, and promoting the common good corresponds to what we are normatively expecting of them.

Still, I agree with Calhoun that we operate with a notion of a *responsible person* that includes as an element a disposition to what she understands as responsibility taking. Even if I find it doubtful that we predictively expect, by default, that individuals do good beyond what is normatively expected, we naturally think of those who do so as responsible persons. In thinking this, we are not necessarily thinking of them as paragons of virtue, or as satisfying *very* high standards: going beyond what is normatively expected is fairly common, if not predictively expected by default. Rather, we take such people to display more of what we see in people who satisfy normative expectations: responsibility takers display *more responsibility* than the compliance responsible, who merely satisfy normative expectations with reasonable reliability.

To accommodate this, I suggest that we think of *responsibility* as it figures in these thoughts as a property or dimension, of which one can instantiate more or less. At a first pass we can think of the relevant property as that of being responsive to what is important, or to normative reasons. The compliance responsible person is indeed responsible, and the person who goes beyond compliance more so. But attributions of responsibility, understood in this way, do not just target persons or agents. They also target behavior: we say that people and institutional agents behave responsibly, or in a responsible manner, and this is naturally understood as saying that they behave in a manner responsive to what is important.

If this is correct, it is natural to think of “responsible” and “responsibly,” as they figure in these contexts, as a gradable adjective and adverb, respectively, or as “gradables,” for short. Gradables signify a property or dimension of which there can be more or less. “Tall,” “wealthy,” “quickly,” and “sensibly” are all examples, as objects can be more or less tall or wealthy and things done more or less quickly or sensibly, instantiating more or less of the relevant dimension. Likewise, someone or something can be more or less responsible and/or behave more or less responsibly, instantiating more or less responsibility.

^{77.} Or to a somewhat strengthened version of the second, where compliance responsibility is understood in terms not just of *basic* normative expectations, however those are understood, but of normative expectations more broadly.

Gradables are often used non-comparatively, as when we say that someone is tall, wealthy, or responsible, rather than saying that they are *more*, or *equally*, or *less* tall, wealthy, or responsible than someone else. What degree of tallness, wealth, or responsibility do we attribute on such occasions? That depends on what the relevant standard is in that particular context:⁷⁸ to say or deny that Jill is tall attributes a different degree of tallness when her kindergarten friends provide the salient comparison class than when we are looking for someone to get a bowl from the top shelf. To say that someone is wealthy might similarly attribute different degrees of wealth when discussing who might be able to afford a good-sized apartment in central London than when discussing global economic disparities. Analogously, I suggest, to say that someone, or some behavior, is responsible is to attribute different amounts of responsibility depending on context.

What standards for degrees of responsibility might be relevant in different contexts? In some context, normative and predictive expectations might set the standard, as can what is required for sharing social practices. In declaring that I will assume that John is a responsible person until shown otherwise, I might plausibly convey that I will assume that he conforms to normative expectations as well as can generally be predictively expected of people: not perfectly, but well enough for whatever social practices we share.⁷⁹ By contrast, in saying that, unlike most of us, Jill handles crises responsibly, and further commending her for being a responsible person, I instead plausibly convey that she instantiates an amount of responsibility that goes beyond what is normally predictively expected, or is expected during a crisis. Underlying this variation, though, is a shared dimension of responsiveness to what is important, which the compliance responsible person and the responsibility taker instantiate to different degrees.

6.5 What Can We Learn from Positive Reactive Attitudes?

Calhoun argues that positive reactive attitudes are revelatory of the responsibility-taking aspect of the default status “responsible person.” I have suggested that, while the default status involves no such aspect, responsibility taking (in Calhoun’s sense) does indeed exemplify a person’s responsibility. Calhoun is also clearly right that responsibility taking is the target of attitudes such as gratitude and appreciation, and I take her to be right that they are significantly different from the negative reactive attitudes. In particular, where resentment and indignation have clear communicative

⁷⁸ The sort of contextualist analysis of gradable adjectives offered here is not uncontroversial, but it is fairly standard and similar points can be made in other semantic frameworks.

⁷⁹ In this, I would be tracking the interest that Calhoun takes to hold the notion of a responsible person in place: that of sharing social practices with others (72–73).

tendencies, prompting demands of recognition of wrongdoing on the part of their targets, attitudes of gratitude or appreciation seem quite different.

Still, I disagree with some of what she says about these attitudes, and about what they can teach us about responsibility. First, it seems clear that gratitude and appreciation are fitting in many cases that do not involve responsibility taking in Calhoun's sense. Gratitude is fittingly directed at the person who jumped into the ice-cold water and saved you from drowning, even if saving you was their duty and not saving you would have been terribly wrong. Moreover, what is fitting is not just gratitude *that* you were saved, but gratitude directed towards your benefactor. What seem to ground gratitude here are the agential resources invested in benefiting you, rather than any supererogatory element.⁸⁰ I take this lesson to extend to numerous, much less dramatic cases of helping that Calhoun describes as instances of responsibility taking, but which I take to be instances of responding to requiring reasons and the importance of certain kinds of comparative weight-giving. Here, too, what grounds the fittingness of gratitude and appreciation is not that these actions go beyond what is normatively expectable, but that they involve investing resources, or taking on costs to help others or contribute to the common good.⁸¹

My second reservation concerns what positive attitudes reveal about responsibility. The negative reactive attitudes have a *structure* that seems to tell us something about what is presupposed by their targets: on Calhoun's appealing and broadly Strawsonian picture, they incipiently communicate demands that targets live up to normative expectations and that they respond appropriately to their failure to do so, both in action and in self-directed attitudes like guilt. Given that demands make sense only when their targets have the capacity to live up to them, the negative reactive attitudes seem to presuppose a range of capacities, both of self-control and of moral cognition and emotion. Given the natural thought that to be responsible is to be fittingly held responsible, and the Strawsonian idea that being targeted by reactive attitudes is the paradigmatic form of being held responsible, the negative reactive attitudes provide a guide to responsibility.⁸²

^{80.} For discussion, see Massoud 2016.

^{81.} I speculate that Calhoun's assumption that gratitude requires going beyond what is normatively expected is part of why she thinks that we predictively expect others to go beyond what is normatively expected. After all, we do predictively expect others to sometimes do things for which it is appropriate to feel gratitude

^{82.} In Michael McKenna's (2012) development of this sort of account, practices of holding responsible are akin to conversations, and the relevant capacities involve capacities to understand the "agent meaning" of actions.

Not all attitudes are structured in rich enough ways to provide such guidance, though. Consider desires. We might standardly desire that people behave responsibly and desire to be treated as responsible persons, but because *promotion of its content* is the only general action tendency of desire, such desires tell us nothing about responsible behavior, or about being treated as a responsible person beyond the idea that these might be things *to be promoted*. There is nothing here that corresponds to the rich interpersonal action tendencies of resentment and indignation. (If we fully understand what it is to desire that people behave responsibly, then we might plausibly understand what it is to behave responsibly. But this is because understanding the *content* of the desire already involves this understanding; the *attitude of desire* towards that content tells us little in addition).

I now suggest that, taken in their full generality, positive attitudes such as praise, admiration, appreciation, and gratitude are like desire in this regard. While we can appreciate and praise someone for their responsibility taking, and be grateful for it, the general action tendencies of these attitudes are not particularly informative. I can praise the weather in Arizona, admire someone's beautiful face, appreciate a fine wine, and be grateful for having been born during times of peace without this involving any recognition of responsibility. At most, certain *forms* of these attitudes might be responsibility recognizing not only in sometimes taking responsible behavior or persons as their content, but in coming with action tendencies that reveal something about their targets.

What forms of positive reactive attitudes are informative? One form in particular seems to provide guidance: the sort of gratitude or appreciation that involves dispositions towards increased goodwill towards their targets. For simplicity, I will now use "gratitude" for this form, and will include goodwill directed towards someone in virtue of their sacrifice not only for the sake of the person displaying the goodwill, but also for the sake of third parties.⁸³ Now, the assumption that something is the fitting target of goodwill does not tell us much about the target, beyond the fact that it has interests. What is revelatory, though, are the conditions to which this goodwill is sensitive. As I previously suggested, we are subject to a variety of balancing norms, telling individuals and groups to give a certain comparative weight to persons and other values over time. I now further suggest that the *goodwill of gratitude* is shaped by such norms, and the importance of a certain balance in weight-giving. The basic underlying explanation of increased goodwill is that (i) the target has given others or the common good weight, and (ii) balance requires that the

⁸³ Eskens (forthcoming) argues extensively for the recognition of this sort of impersonal gratitude.

target is now given more weight than would otherwise have been the case: acting in their interest or on their preferences has now become more important.⁸⁴

Here I can only briefly sketch the kinds of balancing norms that I take to be at work. In the simplest case, where one person gives weight to the interests of another and the benefactor is grateful, a central balancing norm is that of reciprocity, which in the case of two equals says that, taken together, the two of them should give each other the same weight over time: if A gives B more weight (compared to A) than B gives to A (compared to B), it becomes increasingly important that they give B more weight compared to A. In ordinary reciprocal relationships, a rough balance is preserved, but in paradigmatic illustrations of strong gratitude, one has done considerably more for the other than vice versa. In cases involving third-party gratitude, I take the relevant balancing norm to be one governing the weight-giving of a group. Society, or even the moral community, might be required to give a certain comparative weight to its members. When one member has sacrificed themselves for another or for the common good, the group has thereby, other things being equal, given less weight to the benefactor than to the individual or collective beneficiary, and it becomes important for the group to give the benefactor more comparative weight. Other members of the group can individually or jointly ensure that they do.

If this is correct, as a rough outline, the kind of goodwill involved in gratitude is structured in a way that tracks the target's role as a weight-giver, as someone subject to weight-giving norms, and as someone who has actually responded to what is important. As this requires that the target is responsive to what is important, it involves recognizing them as a responsible person. For this specific form of gratitude, then, I agree with Calhoun that it is interestingly responsibility recognizing. Importantly, it is responsibility recognizing in a way that mirrors the responsibility-recognizing character of resentment and indignation. Those attitudes, in their responsibility-recognizing forms, characteristically involve withheld goodwill, on the grounds that their targets have given more comparative weight to their own interests or judgments than morality allows, and have given too little weight to individual victims or the common good. However, to correct the resulting imbalance, fitting targets of resentment and indignation will have to give their own interests and point of view less comparative weight, as characteristic of the humbling stance of sincere apology for culpable wrongdoing and accompanying actions to repair what has been

⁸⁴ In saying that the basic underlying explanation of goodwill is prior weight-giving, I am not assuming that all weight-giving matters for the relevant balance. Even if we can deserve gratitude for saving someone's life when it is our duty to do so, we do not in general deserve gratitude for weight given to someone against their will or in doing what we are normatively expected to do and that others rely on as a matter of course. See footnote 75.

harmed.⁸⁵ The need for these responses on the part of perpetrators explains why the negative reactive attitudes involve significant action tendencies directed at specific uptake from their targets in a way that gratitude does not.⁸⁶

6.6 Concluding Remarks

We do well to follow Calhoun in considering a wider range of responsibility-related phenomena, and to consider the often-neglected positive aspects of responsibility. In this commentary, I have followed her example, and followed her to some of her conclusions. We do indeed predictively expect people to help others and promote the common good in ways that merit gratitude, and when people satisfy these expectations, they are indeed acting responsibly, and being responsible persons. Moreover, the demands of morality leave us with significant freedom in deciding when and how to contribute to the common good. But I have parted ways with her in suggesting that the good behavior predictively expected of others, by default, is also normatively expected, and explained how that is compatible with our sense that we are often free not to help or contribute to the common good, but merit gratitude when we do. In addition, I have argued that what the positive reactive attitudes reveal about their targets is not primarily that they are responsibility takers, but that they are weight-givers, subject to balancing norms. In this regard, they mirror the negative reactive attitudes.

I take this departure from some of Calhoun's conclusions to be based on the broadly Strawsonian methodology that she is following, using attention to our practices of holding responsible as a guide to what it is to be responsible. Importantly, these practices are not free-floating, but in turn responsive to the structure of our normative expectations. Appealing broadly to expectations of proper regard that are central to interpersonal life, Strawson (1962) explained why resentment and indignation would be undermined by standard excuses and exemptions. Here I have appealed to a more detailed characterization of expectations at work, as revealed by reactions to failures to help on particular occasions and failures to contribute sufficiently to the common good over time. What I have suggested is that such reactions cast a different light than the one offered by Calhoun on our expectations that people will contribute to the common good and our reactions of gratitude when they do. Gratitude can be fitting even for what was normatively expected, and what we are predictively expected to do—to give sufficient weight to the common good

^{85.} For discussion, see Björnsson 2022.

^{86.} I take it that there is no corresponding requirement on targets of gratitude to give more weight to themselves compared to their beneficiaries in the future. We may sacrifice our own interests in a way that we may not sacrifice the interests of others.


over time—is also normatively expected, as balancing norms are a pervasive part of the normative landscape.

Acknowledgments

Many of the ideas I try out here have been sharpened in ongoing collaboration with Romy Eskens on related issues. The chapter has also benefitted from comments from audiences at the Descartes Lectures at Tilburg University, the Higher Seminar in Practical Philosophy at Stockholm University, and the Practical Philosophy and Political Theory Seminar at the University of Gothenburg, as well as from comments by Miguel Egler and Alfred Archer. All this help is greatly appreciated. Finally, I want to thank Cheshire Calhoun for bringing to the fore a wealth of new questions and stimulating ideas, for highlighting the importance for responsibility of phenomena that the literature has mostly left out of sight, and for giving me this opportunity to comment on her work.

6.7 References

- Björnsson, Gunnar. 2022. "Blame, Deserved Guilt, and Harms to Standing." In *Self-Blame and Moral Responsibility*, edited by Andreas Brekke Carlsson, 198–216. Cambridge: Cambridge University Press.
- Björnsson, Gunnar and Bengt Brülde. 2017. "Normative Responsibilities: Structure and Sources." In *Parental Responsibility in the Context of Neuroscience and Genetics*, edited by Kristien Hens, Daniela Cutas, and Dorothee Horstkötter, 13–33. Cham: Springer International Publishing.
- Bykvist, Krister. 2022. "Wellbeing and Changing Attitudes across Time." *Ethical Theory and Moral Practice*.
- Eskens, Romy. Forthcoming. "Moral Gratitude." *Journal of Applied Philosophy*.
- Garfinkel, Harold. 1964. "Studies of the Routine Grounds of Everyday Activities." *Social Problems* 11: 225–50.
- Hill, Thomas E., Jr. 1971. "Kant on Imperfect Duty and Supererogation." *Kant-Studien* 62: 55–76.
- Korsgaard, Christine M. 1996. *The Sources of Normativity*. Cambridge: Cambridge University Press.
- Massoud, Amy. 2016. "Moral Worth and Supererogation." *Ethics* 126: 690–710.
- McGeer, Victoria. 2019. "Scaffolding Agency: A Proleptic Account of the Reactive Attitudes." *European Journal of Philosophy* 27: 301–23.
- McKenna, Michael. 2012. *Conversation and Responsibility*. New York: Oxford University Press.
- Pummer, Theron. 2023. *The Rules of Rescue: Cost, Distance, and Effective Altruism*. New York: Oxford University Press.
- Shoemaker, David. Forthcoming. *The Architecture of Blame and Praise*. Oxford University Press.
- Stohr, Karen. 2011. "Kantian Beneficence and the Problem of Obligatory Aid." *Journal of Moral Philosophy* 8: 45–67.
- Strawson, Peter F. 1962. "Freedom and Resentment." *Proceedings of the British Academy* 48: 187–211.
- Williams, Bernard. 1981. "Persons, Character, and Morality." In *Moral Luck: Philosophical Papers 1973–1980*, edited by Bernard Williams, 1–19. New York: Cambridge University Press.



REPLIES TO COMMENTATORS


My thanks to Gunnar Björnsson, Jules Holroyd, and Heidi Maibom for this extraordinary set of comments. I particularly appreciate the way that each, in engaging with my lectures, lays out their own substantive, thought-provoking philosophical project. The result is three essays that are philosophically illuminating in their own right, and not just as insightful and challenging commentaries. In different ways they offer important challenges, amendments, and extensions to central features of my account of responsible persons.

Gunnar Björnsson focuses on the *three*-dimensionality of my account, arguing that we don't need a default presumption that responsible persons are responsibility takers. We don't need that presumption because much of what I take to be non-obligatory and elective promotion of goods within practices (and thus the work of responsibility-takers) is in fact normatively expected. If that's right, then the only basic presumptions about responsible persons that are needed to explain the structure of social practices are that they are accountable for failures to live up to normative expectations, and possibly also that they are compliance responsible.

Jules Holroyd and Heidi Maibom draw attention to my apparent idealization of social practices, especially the norms involved in those practices. They agree that in hierarchical societies social norms typically protect the interests of the privileged and powerful and serve to maintain social hierarchy. Holroyd not only suggests that greater attention to the non-ideality of actual social practices, especially the way reactive attitudes function, is in order. More critically, she proposes that a social practice account of responsible persons should *begin* from descriptions of how responsibility practices work in hierarchical societies. What is at stake here is the appropriate methodology for constructing a social practice account of responsible persons.

Heidi Maibom, after noting that social morality and practices of responsibility have historically served the interests of the privileged, moves in a different direction from Holroyd. My discussion of compliance responsibility raises a more fundamental question about how social cooperation gets off the ground. She argues that it cannot get off the ground if cooperation depends on being able to predict how potential

cooperators will behave. Rather, it is established roles, practices, and norms that play the explanatory role, not predictions. While her comments are largely compatible with my own account of compliance responsibility, she is skeptical about my talk about normative-cum-predictive expectations. Talk about predictive expectations naturally suggests that individuals' compliance with social norms depends on a capacity to predict what others will do. She argues that this is not the case.



7

BJÖRNSSON ON TAKING RESPONSIBILITY

7.1 Preliminaries

Gunnar Björnsson and I disagree most centrally on the question of what is normatively expectable: Do we normatively expect people to make contributions to the common good and others' welfare? Or do we hope that, as responsibility takers, they will elect to make such contributions, contributions that are not normatively expectable? Since there is some risk of misunderstanding what we are disagreeing about, let me begin by clarifying what “normative expectation” means here before turning to my response. The term “normative expectation” does not have a univocal meaning. If one is doing ethical theory, “normative expectation” refers to *moral* expectations that are genuine, correct, rationally justified expectations—what Heidi Maibom describes as “Platonic” morality, or what Jules Holroyd might describe as constituting the ideal morality system. I am not doing ethical theory or presupposing any ethical theoretical view. I use “normative expectation” to refer to whatever is socially normative in a particular social world. These expectations may not be morally normative in that social world, since many practice norms, while social norms, are not social *moral* norms. A workplace, for example, may have normative expectations about professional attire that are not socially construed as moral expectations. Thus, when I ask what conception of responsible persons is presupposed by the structure of our social practices, I am not just concerned with *moral* responsibility.

Both Maibom and Holroyd quite nicely draw attention to the fact that social moral norms may diverge widely from what we as ethical theorists, or as critically, reflective individuals, think the genuine, correct, rationally justified moral norms are. As I use the term, “normative expectation” applies to *whatever* the normative expectations are in a particular society. Given this, the enforcement of social norms via blame, and, as Holroyd points out, also via praise, can be oppressive and disadvantageous to many people if those norms are themselves oppressive and

disadvantaging. Similarly, “practice good” also applies to *whatever* is taken to be the good underpinning a practice. A practice might be organized around the good of racial purity, as marriage practices once were by law in the U.S. That something is a practice good does not mean it is good from a critical normative point of view. That individuals elect to promote a practice good in non-required ways does not make what they do genuinely good from a critical, reflective point of view. Thus, being compliance responsible in a particular society may mean complying with wrongly oppressive norms. And being a responsibility-taker may mean promoting morally objectionable goods as viewed from a critical, reflective point of view.

7.2 Engaging with Björnsson

Gunnar Björnsson’s description of what is normatively expectable looks a great deal like, even if it is not intended to be, the sort of thing that an ethical theorist would say. It may be true in “our” social world that “people are normatively expected to be willing to contribute to the common good and help others to some extent” (Björnsson, page 4 of this volume); but this needs to be taken as an empirical claim about “our” norms in our social world—as do any claims I make about norms in “our” social world. This is to say that Björnsson and my disagreement about what is normatively expectable needs to be read as a disagreement about how best to describe the social moral norms operant in “our”—generally, Western—social worlds. It is not a disagreement about the genuine, correct, rationally justifiable normative expectation.

What Björnsson and I fundamentally disagree on is how to characterize our shared (broadly, Western) social world. Is a conception of persons as responsibility-takers presupposed by the structure of *our* social practices? I think so. He thinks not. I will shortly return to this disagreement. But let me begin with the central challenge that provides the context for this disagreement.

The central challenge Björnsson poses might be put this way: *Given that many of our social practices appear to be norm-and-election structured, does it follow that our practices rely on a default assumption that participants are responsibility takers, i.e., disposed on at least some occasions to elect to promote practice goods in non-required ways?* In short, is being a responsibility-taker part of the status “responsible person”?

As I understand him, there are two central reasons for a negative answer. First, we don’t need the third responsibility-taking dimension to make sense of the norm-and-election structuring of social practices. That is, we don’t need a *distributive* default

assumption that *all* social participants, absent disqualifying evidence, are not only *capable* of understanding why some non-required actions would be good to do, but are also at least sometimes *disposed* to elect good but non-required actions.

To explain the norm-and-election structure of social practices, all we need suppose is that social participants have the capacities relevant to accountability and compliance responsibility. Consider the capacity for reasons-responsiveness that is constitutive of being accountability responsible. Responsiveness to reasons includes not only responsiveness to the reasons favoring what is *normatively expected*. It also includes a capacity to respond to reasons favoring those *good things that a practice treats as elective*. The socialization undergirding compliance responsibility ensures that social participants have a *realized* capacity to respond to reasons. But responsiveness to reasons, Björnsson rightly observes, comes in degrees: some people are more responsible than others in the sense of being more responsive to reasons—more responsive to what is important. The word “responsible” thus can be used to pick out a gradable property. The most responsible people are those who elect to do non-required good things. Responsibility-taking is thus an *achievement* deserving *appraisal* respect (to use Stephen Darwall’s (1977) contrast between appraisal and recognition respect). “In commending someone for being responsible, one conveys that she instantiates an amount of responsibility that goes beyond what is normally predictively expected” (16).

Thus, to explain the norm-and-election structure of social practices, we need not assume that responsibility-taking is part of a default *status* that deserves *recognition* respect. We need only assume that in a sufficiently large group—perhaps the whole society—*enough* individuals can be predictively expected to make non-required elections. “Given natural variation in people’s dispositions to contribute beyond what is normatively expected [this] non-distributive prediction is no doubt correct for large enough groups of participants” (5).

Björnsson’s explanation of practices’ norm-and-election structure seems simpler and more modest than my explanation. Less is imputed by default to social participants, and responsibility has two rather than three dimensions.

Björnsson is surely correct that we use “responsible” as a grading property. Some people are more responsible than others. And it’s surely material to a particular practice *how many* people will be interested in promoting *that* practice’s good. For example, if you have an educational institution, and you think that very few teachers will bother to continuously improve their teaching, then it makes sense to

make additional pedagogical training required. If, on the other hand, you think many teachers will, on their own initiative, work to continuously improve their teaching, then it makes sense to leave additional pedagogical training up to individual election.

Were this all to Björnsson's view, it would lose something important. In regarding responsibility-taking as *only* an achievement of some, what we lose is a stance toward people in general of presuming a disposition to take responsibility for promoting goods on at least some occasions. One implication of his achievement view is that it's appropriate to ferret out the responsibility-achievers and direct your requests for volunteers, for donations, for help, for favors only at them. Besides overburdening the highly responsible few, I think this conveys a disrespectful message to those not tapped: "You're not responsible enough for us to bother asking you."

But let's turn now to Björnsson's second and more central reason for thinking that we don't need the third responsibility-taking dimension to make sense of the norm-and-election structuring of social practices. Taking some (perhaps a lot of) interpretive license, I understand his suggestion to be this: although it may *seem* that social practices have a norm-and-election structure, much of what appears elective is not in fact.

Consider teaching. Teachers are required to show up for their classes and provide something of educational value. Failure to do so may be met not just with resentment, but more punitive responses. Beyond these strict requirements, many educational activities are left to teacher discretion—provision of classroom resources, amount of time spent mentoring individual students, taking courses to improve pedagogy, and the like. As I interpret Björnsson's view, the fact that a *particular* action is not required does not mean that there is no sense in which it is normatively expected.

Björnsson proposes normative expectation: "People are normatively expected to be willing to contribute to the common good and to help others to some extent, giving the ends of others and shared practices some weight" (5). Our resentment of those who never contribute or help despite plentiful opportunities to do so is evidence that this is indeed our normative expectation of participants. When, how, and how often to act on this normative expectation will necessarily, on his account, involve individual—we might say, discretionary—judgment. This is for two reasons. First, promoting common goods and others' interests comes into competition with the individual's own activities. Those activities have various degrees of importance for the individual. Individuals are the best, or the only, ones positioned to say just how important an activity is for them. Thus, they must be the ones to decide whether an opportunity to

promote the common good or others' interests is outweighed by the importance of their own, conflicting activities. Individuals have "evaluative autonomy" in this respect (9). Second, "it is important that we give others and the common good a certain weight over time compared to the weight we give ourselves" (11). Again, the individual is the one best positioned (or the only one positioned) to assess whether, over time, they have complied with what Björnsson calls "balancing norms." Were our practices to strictly require everything that it would be good to do, this would interfere with agents' efforts to balance their own interests against others' and to do so over time. Hence the importance of structuring practices around both strict requirements and elections. In short, what individuals "elect" to do isn't typically normatively optional or supererogatory.


I find Björnsson's proposal attractive. It plausibly explains the norm-and-election structure of practices. Some version of what he calls normative expectation does indeed seem to describe our social world. And recognizing that *all* are under the normative expectation to contribute to the common good and help others to some extent removes my earlier worry about targeting responsibility achievers for requests. A *stance* toward *people in general* of expecting they will be willing to volunteer, donate, provide, help, do favors is warranted because there's a *normative expectation* of accountability/compliance responsible persons that they will promote common goods and others' interests on at least some occasions. The good-doing I called common decency is a *basic* normative expectation.

I said at the beginning that, to my mind, what Björnsson and I fundamentally disagree on is how to characterize our shared (broadly, Western) social world. Björnsson claims that in our social world we are always under a normative expectation to contribute to the common good and others' interests *unless* our own activities at the time are more important to us or *unless* over time we've exceeded the balancing norms and given more weight to others and the common good than required. This (in his words) "extensive" (7, 9) normative expectation sounds to me like what a moral philosopher would (and quite likely should) say we are genuinely required to live up to. I'm dubious that it describes a shared moral understanding in our own social world.

Björnsson's claims that we have strong normative expectations that others will deliver *minor* forms of help, and that we resent those who refuse to do *anything*, do strike me as true of our social world. Many minor forms of help are matters of common decency. I take norms of common decency to be basic norms that any minimally well-socialized individual can be expected to comply with. So we don't

disagree here. But these facts do not entail the further claim that we accept the extensive normative expectation he describes. That expectation is extensive in the sense that it is always in play unless an exempting justification is available. Those exemptions are either that the individual's own activities are more important or that the individual has satisfied the balancing norm for weighting others' interests versus their own over time. What exactly the normative expectation is in our social world for contributing to the common good and promoting others' interests is, of course, an empirical question. It is also an empirical question whether our social practices operate on the assumption that one must always be able to justify not promoting the common good and others' interests. I'm dubious that our own social world is governed by this extensive normative expectation, deviations from which must be justified. For my own part, I find Peter Singer's (1972) observation that we think of charity, even for the most desperate, as supererogatory closer to the empirical truth. My own, depressing sense is that we live in a social world in which there is *some* expectation of contribution to the common good and others' interests but also a very large zone in which individuals need not justify their non-contribution by citing legitimate countervailing reasons. It's for that reason that I think the norm-and-election structure of our social practices involves genuine elections—responsibility taking.⁸⁷ If so, the three-dimensional model of responsible persons that I propose applies to our social world. That said, it's important to keep in mind that my goal was not to describe our specific social world. Rather, my aim was to develop an account of responsible persons that applies to social worlds that have varying structures: they might simply have norm-structured practices; they might in addition have some or many decently functioning social practices; and they might have genuinely norm-and-election-structured practices. Any social world in which Björnsson's normative expectation does not operate, and which has norm-and-election structured practices, will need to presume that participants are responsibility takers in my sense.

⁸⁷. There are, however, two ways of interpreting the normative expectation principle that bring it more into line with social reality and less in line with an ethical theorist's view. One is to emphasize how unconstrained judgments of personal importance are: individuals have wide latitude—"evaluative autonomy"—to heavily weight the importance of their own (their kin's, their social group's, etc.) activities. The other is to acknowledge that, if the aim is to describe *our* social world, the balancing norms can't be specified by philosophers, but must be socially constructed (14). It is an empirical question what those balancing norms are.



8

HOLROYD AND NON-IDEALIZING ACCOUNTS

8.1 Preliminaries

As I indicated at the beginning, I take it that Jules Holroyd and I disagree—or at least appear to disagree—on the methodology for constructing a social practice account of responsible persons. Since I do emphasize beginning from “the facts as we know them,” and I do use many examples of our own social practices, there’s some risk of misconstruing what I take a social practice account to involve. In particular, there’s some risk of taking my central aim to be one of describing in detail what responsibility looks like in our current social world. So let me try to clarify my understanding of what constructing a social practice account amounts to.

I’ll begin by noting that, unlike Strawson, I don’t think the account should begin from facts about our *practices of responsibility*, including reactive attitudes, since, as I argue in the lectures, not all of our practices of responsibility and responsibility-recognizing attitudes may be salient. The social practice account of responsible persons that I offer thus begins not from practices of *responsibility*, but from the vast array of cooperative social practices oriented around some good—practices of parenting, hosting parties, sharing trains, working in jobs, and so on. The central questions, then, are:

1. Given that social practices are *norm-structured*, what *developable* if not realized capacities must participants be presumed to have if there are to be any *norm-structured* practices at all?
2. Given that a *well-functioning* (or, as I sometimes say, *decently functioning*) social practice is one that doesn’t depend for its functioning on continuous surveillance and coercive enforcement, what *developed* capacities must participants in those well-functioning social practices be presumed to have?

3. Given that some (perhaps many) social practices are *norm-and-election structured*, what capacities must participants be presumed to have in order for such a structure to make sense?

Because I draw so heavily on what I take to be the “facts as we know them” about “our” social world, an important feature of the account may drop from view: it is intended to be useful in understanding the nature of responsible persons in social worlds with quite differently structured social practices. For example, a social world might have no, or only some, *decently functioning* social practices: that world relies entirely or in part on surveillance and coercion to assure compliance with social norms. In such worlds (or parts of worlds), responsible persons will at best be presumed to be accountability responsible, not compliance responsible or responsibility takers. Conceivably there are highly demanding social worlds in which all that is good to do is also normatively expected. In such worlds, there would be no room to conceive of responsible persons as responsibility takers (a point that I take Björnsson to be making). There are also conceivably—although this seems quite hypothetical—social worlds where participants are so strongly disposed to promote practice goods that there is no social point to having enforceable norms. Their practices would not be norm-structured, and their conception of responsible persons would not include either accountability or compliance responsibility.⁸⁸

What I want to stress is that I focus on “our” social world as a *useful device* in illuminating the essential features of (1) all social practices that we know of (they are norm-structured), (2) decently functioning social practices (they do not rely on monitoring and coercive sanctions), and (3) a prevalent (or at least possible) type of social practice (it is norm-and-election structured). I then work out what *must be presupposed about social actors if there are to be practices of those three kinds*. My claim that there are *default* presumptions of responsibility tied to any social practice, to decently functioning ones, and to norm-and-election structured ones is thus not an empirical claim.

If one does not keep these points in mind, it will be easy to overread my claims about a *default* presumption. That there is a default presumption that individuals have the developable or realized capacities connected with accountability is a claim about *participants* in a practice. Not everyone who is *involved* in a social practice may be regarded as a social participant. Infants and extremely young children are involved in the social practice of the family before they have even a developable

⁸⁸. Kant suggested that the notion of an imperative—a demandable normative expectation or requirement—does not apply to God. So, one might imagine a society of saints or gods.

capacity to understand norms. Enslaved Africans under chattel slavery were involved in the social practice of farming, but largely as tools to be managed rather than as participants. That there is a default presumption that individuals are compliance responsible is a claim only about participants in *well-functioning* social practices. My claim is not that the default presumption of compliance holds within all possible social practices regardless of whether they are well-functioning or not. The extent to which a presumption of compliance is in place gives us a measure of the degree to which a practice is well-functioning.

8.2 Engaging with Holroyd

Holroyd makes three central claims: (1) I work from a “rather idealized conception of social practices” (5); (2) it is important to pay “attention to more non-idealized practices” in our and other societies’ hierarchically structured social world (2); and (3) in working out an account of responsible persons, “we have good reason to *start* with the non-ideal forms of social practices we find, and the ways in which social hierarchy and power dynamics inflect our responsibility” (19).

Starting with the first point: do I have an idealized picture of social practices? My examples, drawing as they do on examples of well-functioning social practices whose practice norms appear acceptable from a critical, reflective point of view, do indeed invite an unduly rosy picture of the empirical reality of what social practices are like in our and similar societies. On this rosy picture, social status does not systematically affect who is accorded full status as a responsible person: all are respectfully and equally extended the default presumption that they have the capacities that make them fit for accountability responsibility within norm-structured practices; all are offered default trust that they will generally be compliant with basic practice norms; and all are presumed by default, and equally so, to be disposed on at least some occasions to elect to promote practice goods in non-required ways. Furthermore, the norms defining particular practices are rosily assumed to be good ones, and thus people are held to account only for what they should be held to account for; the values that practices serve are good ones, and thus responsibility-taking advances something genuinely worthwhile.

Unless the non-ideal is called into view, it is all too easy for talk about “normative expectations” to get covertly replaced with “*reasonable* or *justified* normative expectations”; talk about being “held to account” to get covertly replaced with “reasonably held to account” (because, after all, the expectations were themselves reasonable); and for my talk about the status “responsible person” being a valuable one to mean that everyone is accorded this status and gets the full value from having

it. As a result, the more favorably positioned “we” does not stop to ask, “For whom are our actual social practices of responsibility working, and for whom are they not?” Both Holroyd’s and Maibom’s cautionary remarks about social reality are perfectly in order.

That said, I do not work from an idealized empirical description of our social practices. As I’ve indicated in the Preliminaries to my engagement with Björnsson, I do not make use of an idealized conception of “normative expectation” that equates “normative expectation” with what, from a critical, reflective point of view, we would regard as correct or justified normative expectations. And, as I’ve clarified in this section, my claims about default presumptions are conceptual claims, not empirical descriptions.

I entirely agree with Holroyd on their empirical points. As they correctly observe, whole social groups might not receive this default presumption of compliance responsibility. Indeed, it is a central feature of the oppression of social groups in hierarchically structured societies that they are distrusted—hence, for example, the over-policing and over-incarceration of Blacks, especially Black men, in our society. This lack of default trust is, as they say, not merely incidental, but systematic. So, as an empirical matter, it is false that all individuals in our society are presumed by default to be compliance responsible. Thus, “whilst the capacities that underpin this status might travel, being recognized as having such a status is not context independent” (6). Importantly, as Holroyd observes, the problem in hierarchical societies is not just that the subordinated are often not trusted to comply (absent threat of penalty), but that the subordinated learn not to trust the privileged to comply with basic norms forbidding assault, murder, patent unfairness, and the like, a distrust bolstered by the lack or insufficiency of penalties for non-compliance.

Holroyd helpfully draws attention to the possibility that social practices might manage to function because they assume that many social participants (specifically, the socially privileged) can be trusted to be compliance responsible, particularly when interacting with fellow privileged individuals, while simultaneously that some subset of participants (specifically, the socially subordinate) cannot. Some social practices thus have a *hybrid* nature: they depend on a substantial number of participants being compliance responsible, but they also assume that a significant subset of participants is not compliance responsible and requires surveillance and coercive sanctions. So, we might say this: the well-functioning of a social practice is a scalar property. The larger the number of individuals who cannot be assumed, by default, to be compliance responsible, the less well-functioning a practice is.

What Holroyd also draws attention to is the fact that the *same* group may be assumed, across social practices, to be non-compliance-responsible. So, we might say this: in diagnosing the degree of oppression in a society, we need to look at the number of social practices in which disesteemed social groups are not presumed to be compliance responsible as a result of stigmatizing stereotypes.

I turn now to her second main claim: it is important to pay “attention to more non-idealized practices” in our and other societies’ hierarchically structured social world. I agree.

A comprehensive theory of responsibility needs to do two distinct things. First, it needs to give us an account of what having the status “responsible person” *consists in*. It needs to specify the capacities that ground having that status and the entitlement to be socially recognized as having it. It needs to specify what attitudes count as responsibility-recognizing attitudes. And it needs to specify, at the most general level, which practices count as responsibility-recognizing practices. Doing these things makes it possible to identify what counts as a status insult. All of this is a conceptual project. As I’ve noted, it makes limited empirical assumptions.

What Holroyd drives home is that the conceptual project is not enough. It invites, even if it does not entail, a rosy, mistakenly idealized, and ideological picture of the empirical realities of our own society. More importantly, the conceptual project is incomplete *as an account of responsibility*. We need, in addition, an empirically informed account of both which practice norms are in place and of how practices of responsibility and responsibility-recognizing attitudes are deployed.

We need that empirical account for assessment purposes. Do the norms governing social practices encode wrongful social hierarchies? Does, for example, the social practice of parenting wrongly make obligatory for women things that are left elective for men? If practice norms are implicated in oppression, deployment of the negative and positive reactive attitudes to enforce or encourage compliance will also be implicated in oppression.

In addition, we need to assess practices of *responsibility*. Do practices of responsibility themselves encode wrongful social hierarchies? A practice of responsibility will, for example, have some place for accepting excuses and for sanctioning norm violations. But what we need to know is whether the norm-violating behavior of some social groups is more readily excused or less forcefully sanctioned than the norm-violating behavior of other social groups. Consider

Holroyd's example of himpathy for white, middle-class rapists.⁸⁹ We also need to know whether members of some social groups are presumed to lack what it takes to be accountability responsible, or compliance responsible, or to be responsibility-takers. That is, do members of some social groups routinely suffer the insult of not being accorded, or fully accorded, the status "responsible person"?

As Holroyd underscores, "if the value of the status 'responsible person' is to be adequately realized, these aspects of the non-ideal social practice that threaten to undermine its value must be recognized and addressed" (8). "It is when we start our theorizing from a different understanding of 'the facts as we know them,' we see many aspects of our responsibility practices ... that are weaponized for oppressive ends. Accordingly, we should ask about the contexts in which the value can be properly realized" (12).


Philosophical work on responsibility typically executes either the conceptual project or the empirically-grounded assessment project, but not both. So, we get, on the one hand, conceptual accounts of what responsibility *consists in* without any attention to how responsibility plays out in our actual social world. Or, on the other hand, we get *critiques* of our social norms and responsibility practices that simply presuppose, without specifying, accounts of what being a responsible person consists in, what counts as a responsibility-recognizing attitude, and what counts as a responsibility-recognizing practice. There's nothing wrong with a division of labor so long as, in the end, we recognize that both labors matter to a fully adequate account of responsibility.

What, now, about Holroyd's last claim: on working out an account of responsible persons "we have good reason to *start* with the non-ideal forms of social practices we find, and the ways in which social hierarchy and power dynamics inflect our responsibility" (19). If the goal is to develop an account of what the status "responsible person" consists in, we *cannot* start with non-ideal responsibility practices. Starting from the non-ideal would (as I understand it) involve starting from *complaints* as the significant "facts as we know them." One kind of complaint concerns the extant social norms. Critiquing the *social norms* to which some are held to account is highly useful for the purposes of improving the normative quality of social practices, including practices of responsibility. I don't see how it is relevant to deriving an account of what responsibility consists in. The latter asks which capacities a responsible person must have, not what it is normatively

⁸⁹. The notion of "himpathy" was introduced by Kate Manne in *Down Girl: The Logic of Misogyny* (2017).

legitimate to hold persons to account *for* or normatively legitimate to expect their compliance *with*.

A different kind of complaint one might start from, in starting from the non-ideal, is a complaint that members of some social group are wrongfully presumed to be non-responsible, or are wrongfully excused from responsibility (as in the case of himpathy), or are excessively presumed to be responsible (as in Holroyd's example of the Black child who is overly treated as a responsibility-taker). Such complaints, however, are based on some understanding of which capacities being a responsible person depends on. The legitimacy of those complaints depends on whether they presuppose a correct account of the responsibility-grounding capacities. In short, starting from the non-ideal this way assumes what the method is supposed to deliver. So, I don't think it's possible to avoid some modest form of ideal theory. We start by asking what it means to be a responsible person: what capacities must they have? Given that account, we're then in a position to identify the ways in which individuals in our non-ideal world are wrongly treated as not responsible persons.



9

MAIBOM AND SOCIAL COORDINATION

9.1 Preliminaries

I find it interesting that my account of the status “responsible person” provoked two of my commentators to critique social norms, and the practices of responsibility that enforce them, by underscoring the connection between privilege and power on the one hand and who benefits or is disadvantaged by social norms, real systems of morality, and practices of responsibility on the other. Heidi Maibom casts her critical eye widely—over the “moral and legal norms that prevailed in Europe during most of its history” (2).

Skepticism about the goodness of morality and worries about the way morality and practices of responsibility “operate just as well as tools of oppression and suffering” (2) is almost never brought to the attention of theorists of *responsibility*. This is surprising when the theorists are Strawsonians. Strawson was, after all, providing one kind of social practice account of responsibility: an account of responsibility drawn from the facts as we know them about holding people to account, and sometimes excusing or exempting them from responsibility. His account would, of course, apply to a hypothetical, ideal social world—a Kingdom of Ends—governed by capital-M morality, in which neither the norms nor the practices of holding accountable and excusing are distorted by systems of power and privilege. But it also holds for our real, non-ideal social world in which we regard each other from the participant attitude. His account of responsibility was, as mine is, neutral between ideal and non-ideal worlds; but that shouldn’t have prevented readers from thinking about what the real-world social practice of morality and of responsibility looks like. Any social practice account of responsibility should provoke reflection on the difference between the two senses of morality, which Heidi observes are not distinguished often enough (12).

In my *Moral Aims: Getting it Right and Practicing Morality with Others* (2016, ch. 1), a central goal was to get those two senses into view and to argue that both senses have a legitimate claim to being what morality is about. The task of normative ethics is to provide guidance in determining *legitimate* normative expectations. The task is to “get it right,” where the “it” is capital-M morality. This is important. It’s what we need if we are to critically assess social norms. But the *social* practice of morality is the only morality game in town. It is only in real social worlds that we share moral understandings, use them to make our actions intelligible to each other, blame, praise, and excuse each other, have normative expectations that are normative-cum-predictive ones, and share understandings of the goods that might be promoted via elective choice to take responsibility. However important it is to adopt a critical, reflective point of view and to try to latch on to capital-M morality (and perhaps also practices of capital-R responsibility), we cannot ignore that morality is also fundamentally a social practice, and responsibility practices and attitudes are fundamentally social practices and socially shaped attitudes.

9.2 Engaging with Maibom

I take Heidi Maibom’s edifying discussion of solutions to the coordination problem as an invitation to me to think more about the relation between compliance responsibility and the coordination problem. The coordination problem is a problem that arises for the individual (person, political party, country, etc.) in contexts where cooperation would be beneficial but where behaving cooperatively with another individual (person, political party, country, etc.) comes with the risk of that cooperation not being reciprocated, of having one’s cooperation exploited, and of being made worse off than if one had not behaved cooperatively in the first place. What should one decide in the face of this problem? Maibom explains why solving the problem cannot depend on our capacity to predict what the others with whom we consider cooperating will do. In short: our and those others’ decisions are too entangled. In reiterated decisions, one might instead adopt the tit-for-tat or tit-for-two-tats strategy, thereby avoiding any need for prediction. But this strategy will not be useful, as she notes, in large-scale societies where we have lots of one-off interactions with others. What she proposes is that social structures, particularly circumscribed social roles, like that of being an art dealer, art purchaser, barista, coffee customer, “take away the need to ascribe psychological states to others for predictive purposes” (6). We simply identify the social context (purchasing art, purchasing coffee), the relevant social roles, and the relevant scripts for those roles (which include norms for executing those roles), and then proceed according to the relevant script. “Instead of honing our skills at psychological prediction, we create predictable people by means of the incentives and disincentives built into the environment” (7).

I, of course, largely agree with Maibom's turn to the central role of social roles and norms. I'm less inclined to describe this as a *solution* to the coordination problem. The coordination *problem* arises for individuals in a context of uncertainty. Where there are established roles, scripts, and norms, this problem is simply *eliminated* so that it never arises in the first place as a "What shall I do?" decision-making problem. The original problem and its social-level solution are a matter of history that predates social structures. Once predictable people have been created by means of the incentives and disincentives built into the environment, there is no longer a problem facing individuals for which they need a strategy.

What might those incentives and disincentives built into the environment be? One is a socially costly system of monitoring and (threat of) sanctions sufficiently dire as to deter the would-be non-cooperator. As I've said, one could have functioning social practices that are, nevertheless, not *decently* functioning ones because they rely on coercive mechanisms to make people predictable. Another is effective socialization so that compliance with at least basic social norms becomes largely automatic. In this case, the familiarity of the social context and what it calls for is the "incentive" built into the environment.

Maibom and I talk as though socialization into norms and roles does pretty much all the work of enabling individuals to be compliance responsible and social practices to be decently functioning. One topic I would have liked to have taken up is the "infrastructure" of responsibility. What environmental structures enable individuals to become, and sustains them in being, responsible persons? Early socialization into existing norms can't do all the work of enabling compliance responsibility. Consider: the point of critiquing existing norms is to replace defective social norms, into which individuals have been socialized, with better ones. For example, we're currently moving toward a new norm of using individuals' preferred pronouns. Ideally, this would become a common decency norm, compliance with which is predictively expectable. It's an empirical question, however, what (broadly) environmental structures would be effective in accomplishing this. What appears not to be effective is the philosopher's preferred strategy of promulgating good reasons. In their study of the requiring norm of female genital cutting in Sudan and the permissive norm of open defecation in Bangladesh, Cristina Bicchieri and Peter McNally (2018) discovered that informational campaigns stressing the reasons for ceasing to do these things (particularly reasons having to do with health) were ineffective. More effective strategies included public figures and celebrities coming out in support of the new norm; changing the evaluatively valenced associations with the old or new behavior (e.g., describing uncut women as "Saleema"—unharmful, pristine, in

a God-given condition); creating conditions for widespread public discussions of the new behavior; the introduction of novel, pejorative terms for those engaging in the old behavior; and telenovelas with relatable characters who conduct their lives according to the new norms.

The infrastructure question isn't just about how to produce compliance responsibility with respect to new norms. It's also a question of how to address social unclarity about what the norms are, or which norms apply in which context. In her spoken comments on the presented lectures, Maibom observed that "the world is a messy place and often situations do not present themselves as calling for application of a relevant norm." So how might the environment enable individuals to determine when to act on a norm they are already disposed to act on?

One infrastructural solution is a meta-norm allowing, or indeed advising, individuals to ask when in doubt. This isn't always in place. I wondered in the Netherlands, for example, if it were gauche to ask my taxi driver whether tipping was expected. Another solution is adequate "signage" about what rules apply in what contexts. That signage might take the form of literal signage—visibly posted rules, such as "No talking in the library," or a poster specifying what to do if one has misgendered someone. Or it might take the form of institutional training, such as workplace sexual harassment training. Or it might take the form of public ethics experts who, like Ms. Manners or Anthony Appiah, publish ethical advice to those who write in with questions.

There are similarly interesting questions to be addressed with respect to the infrastructure of responsibility taking. You might agree with Björnsson that the norm-and-election structure of social practices depends only on there being a sufficient number of individuals who have the gradable property of being responsible. Even on that view, there's reason to think about what the environmental incentives or disincentives to responsibility might be. Widespread transactional approaches to securing volunteers, for example paying blood donors for donating blood, may work as an environmental disincentive to being a responsibility-taker, even if it is an effective strategy for recruiting more volunteers.

Let me turn now to where Maibom and I appear to disagree, although I don't think we actually do. I distinguish the basic normative expectations relevant to compliance responsibility by describing them as not merely normative (concerned with what individuals ought to do) but normative-cum-predictive (concerned with what individuals both ought to do and may be predictively expected to do). Maibom takes

me to be saying that “predictive expectations underlie much actual cooperation” (8) in the sense that the *decision* about whether to cooperate or not—whether to comply with a norm or not—on some specific occasion depends on the person’s predictions about what a potential cooperation partner believes, wants to, intends to, or will do (8).

In describing our normative expectations with respect to basic norms as normative-cum-predictive, I was not aiming to describe anything about an individual’s decision-making process, but rather the *character of the normative expectation*. It is one thing to have a purely normative expectation: I think, “This is how individuals *ought* to behave.” So, I might say of a society in which ethnic cleansing is taking place that individuals ought not to do this. Or I might think that politicians ought not to promulgate falsehoods, even though many regularly do. It is another thing to have a normative-cum-predictive expectation: I think, “This is both how individuals *ought* to behave and in fact how *most* individuals manage to behave.” So, I might say that individuals ought to hold the door open for people immediately behind them, and most people manage to achieve this level of courtesy. Or I might say that people should pay for the items they take out of stores, and most people manage to do so rather than stealing. These are claims not about my own decision-making process but about the nature of the normative expectation. It makes a difference whether my normative expectation is purely normative or normative-cum-predictive. Failures to live up to normative-cum-predictive expectations invite a specific kind of resentment (if the failure concerns how I am treated) or indignation (if the failure concerns how others are treated), namely, *incredulous* resentment or indignation that the individual has failed to do what minimally well-socialized people routinely manage to do. I might experience incredulous indignation at what some newsworthy person (say, a prominent political figure) has said because it so astonishingly violates basic norms of honesty. Recall the parents’ reaction, in Garfinkel’s study, to their children treating them exceptionally formally and politely. The intensity of reactive attitudes to violations of normative-cum-predictive expectations might be especially intense for other reasons as well—for example, they not only involve doing what people ought not to do, but also help to undermine a valuable social norm.⁹⁰


I agree with Maibom when she notes both that “if social norms are in place, and the society is not itself dysfunctional, then one can predict that people will adhere to them,” and that “it is not the prediction itself that underwrites the adherence” (8). It is the sheer fact that it is a (basic) norm. Other motives may also enter into the causal explanation for norm adherence: fear of sanctions, belief that the norm is legitimate and ought to be complied with, giving weight to the fact that others are

⁹⁰ Thanks to Alfred Archer for this point.

counting on one, a desire to make one's behavior intelligible to others, default or non-default trust that others will also comply with practice norms, and so on.

There are, however, two senses in which predictions do underlie adherence to social norms. First, as Maibom observes, predicting particular individuals' intentions is obviated when people in general are predictable. Second, general predictability underlies our assessment of what the norms are. What social norms require, prohibit, or permit changes over time alongside changes in general patterns of behavior of those participating in a social practice. Get enough defection, and you no longer have a social norm. In the 1950s, dressing up for air travel was the "done thing." It no longer is; wear whatever you like. To get a fix on what the current social norms are, including the norms built into the scripts for different social roles, it helps to determine what is predictively expectable of people generally or of people in particular roles.

In sum, I don't think Maibom and I really disagree. We both think that the kind of social cooperation involved in compliance responsibility does not depend on predicting how particular potential cooperators will behave. And we agree that socialization produces predictable people.



10

A SHORT NOTE ON GRATITUDE, PRAISE, AND TRUST

In my lectures, I draw attention to the responsibility-recognizing attitudes connected with each of the three dimensions of responsible persons. For example, I emphasize that basic trust is the principal *compliance-responsibility*-recognizing attitude and that gratitude, appreciation, and praise are the principal *responsibility-taking*-recognizing attitudes. Björnsson and Holroyd both underscore that there is no neat alignment between gratitude and praise and responsibility taking. Björnsson observes that there are obvious cases where a person's doing what they are obligated to do merits gratitude. Fulfilling a duty of rescue at significant cost to the agent is a case in point. "What seems to ground gratitude are the agential resources invested in benefiting you, rather than any supererogatory element" (17). So, gratitude is not limited to instances of responsibility-taking, but is also connected with accountability responsibility. In a similar vein, Holroyd observes that praise can be an effective and apt expressive tool for developing responsibility-relevant capacities" (14). An obvious case is where one thanks and/or praises a child for telling the truth as a way of both reinforcing this norm for them and motivating them to act on it in the future. Here again, positive reactive attitudes are not exclusively aligned with responsibility taking, but can also recognize the individual's status as accountability responsible.

I entirely agree. My own example of this lack of alignment concerned resentment, which in its paradigmatic form recognizes accountability responsibility, but in the form of incredulous resentment recognizes compliance responsibility.

Björnsson and Holroyd both go on to offer their own, richer accounts of gratitude and praise. Björnsson uses his notion of balancing norms to understand the kind of grateful reciprocity called for when an individual helps others or contributes to the common good. Holroyd focuses on developing a deeper account of how praise


functions, particularly the way it communicates the praiser's values and seeks (and sometimes pressures) uptake of those values via accepted praise.

Maibom takes up trust, informing us of some of the fascinating empirical literature that sheds light on both default and non-default forms of trust. I found particularly compelling her reminder of the role that a structured environment plays in trusting strangers. In her interactions with an art gallery owner, she observes, "it wasn't the owner I trusted so much as the very structure within which our interaction took place" (9). It's an excellent reminder that trust in others' compliance responsibility is deeply enmeshed with trust in the larger social structure in which social participants operate.

Let me end by again thanking Gunnar Björnsson, Jules Holroyd, and Heidi Maibom for their extraordinarily thoughtful and highly substantive comments on my lectures. I hope that my responses go some way toward persuading them that there is utility in conceiving of responsible persons, in at least some social worlds, as responsibility-takers; in the necessity of pairing a society-neutral account of responsible persons with a description and critique of how our actual responsibility practices work; and in thinking that there is a useful distinction to be drawn between those normative expectations that are purely normative and those that are normative-cum-predictive.

References

- Bicchieri, Cristina and Peter McNally. 2018. "Shrieking Sirens: Schemata, Scripts, and Social Norms. How Change Occurs." *Social Philosophy & Policy* 35 (1): 23–53.
- Calhoun, Cheshire. 2016. *Moral Aims: Getting it Right and Practicing Morality with Others*. New York: Oxford University Press.
- Darwall, Stephen L. 1977. "Two Kinds of Respect." *Ethics* 88 (1): 36–49.
- Manne, Kate. 2017. *Down Girl: The Logic of Misogyny*. New York: Oxford University Press.
- Singer, Peter. 1972. "Famine, Affluence, and Morality." *Philosophy & Public Affairs* 1 (3): 229–43.



Editors
Alfred Archer
Miguel Egler