

## Group Agency Meets Meta-Ethics:

### How to Craft a More Compelling Form of Normative Relativism

#### I. Introduction

In this paper, I seek to show how we might draw upon the idea of *group agency* in order to offer a more compelling relativist view of moral reasons and other kinds of normative reasons, too.<sup>1</sup> It will be a central tenet of the view I propose that whole societies, as well as individual persons, can be genuine agents and can possess normative reasons in light of their aims. The resulting meta-normative theory shares key advantages in common with previously defended relativist views. Yet it also, I will argue, holds the potential to avoid some of their most problematic consequences.

There are at least two well-known motivations for rejecting moral realism and for endorsing a form of antirealism, such as relativism, instead. By *moral realism*, I mean the view that (i) there are moral facts, which (ii) our moral beliefs and judgments purport to represent, and which (iii) are all fundamentally attitude-independent.<sup>2</sup> I refer to anyone who rejects moral realism as a *moral antirealist*. I use the term *relativism* broadly to include any antirealist view according to which there are moral facts (or normative facts of any kind) that are fundamentally dependent upon contingently held attitudes such as beliefs, desires, values, or aims.<sup>3</sup> This is to

---

<sup>1</sup> Besides moral reasons, my focus here will be reasons of self-interested practical rationality.

<sup>2</sup> I am roughly following Shafer-Landau's (2012) definition. He characterizes the realist's commitment to the attitude-independence of the moral facts as follows; our moral "beliefs, when true, are not true by virtue of being the object of, or being implied by, the attitudes of (even idealized) agents" (1). I include the word "fundamentally" in clause (iii) to stress the following point. A realist might hold that some of the moral facts are attitude-dependent, but that all of these facts are made true at least in part by some further attitude-independent moral fact(s). The antirealist will insist that all normative facts (if she grants there are any) are attitude-dependent.

<sup>3</sup> I mean for this to include forms of cultural moral relativism, as defended e.g. by Harman (1996) and Velleman (2015), as well as Subjectivism about all normative reasons, as defended e.g. by Street (2006) and Sobel (2016). Outside the scope of my discussion here are Kantian Constructivist views, according to which there are some

deny clause (iii) of the definition of realism. The error theorist, on the other hand, denies clause (i), while classic non-cognitivism is distinguished by its denial of clause (ii).<sup>4</sup>

I have alluded to two influential lines of argument for antirealism. I will discuss them only briefly; it is not my intention to mount a defense of them here. The first takes the form of an epistemological challenge for moral realism, or for realism about normativity in general. One way of articulating the worry is that the realist might be unable to explain how it is that we could reliably form true beliefs about the moral (or normative) facts, given their attitude-independence. Assuming that we do reliably form mostly true normative beliefs, one natural solution to these epistemological difficulties is to embrace a theory according to which all of the relevant normative facts are fundamentally attitude-dependent rather than attitude-independent. This is, for instance, Street's (2006) proposal. The hope is that such a theory would make room for a plausible epistemological story that appeals to the dependence of normative facts upon descriptive facts about us and our psychology in order to explain how it is that we can, and often do, discover what those facts are.<sup>5</sup>

A second line of argument, which seeks to provide support especially for cultural moral relativism, proceeds by emphasizing the widespread observable diversity of value systems espoused by people of various historical communities. Harman (1996), for example, claims that

---

normative truths that, instead of depending upon any agents' contingently held attitudes, are made true by features of rational agency as such.

<sup>4</sup> The error theorist holds that there are no moral facts, even though that is exactly what our moral judgments purport to express, and so our moral judgments are systematically in error. See Mackie (1977). According to classical forms of non-cognitivism, on the other hand, our moral judgments are not truth-apt at all; they are more like expressions of feeling or preference than statements of alleged fact. Classic non-cognitivist views include forms of emotivist expressivism as popularized by Ayer (1936). I use the term "classic" to mark a distinction between these views and more recent forms of "quasi-realist" non-cognitivism, as defended e.g. by Blackburn (1993, 2005). What distinguishes quasi-realism is the claim that our commitment to the existence of moral facts (and to their objectivity) is consistent with a non-cognitivist characterization of moral discourse. I will not say anything further here about error theory or non-cognitivism.

<sup>5</sup> For detailed discussion, see also e.g. Bedke (2009, 2014); Berker (2014); Copp (2008); Enoch (2010, 2011); Joyce (2006, 2013, 2016); Kitcher (2005); Schechter (2010); Setiya (2012); Shafer-Landau (2003, 2012). See also my own Dyke (2019).

the following thesis is a “reasonable inference” as a plausible explanation of observable cultural diversity:

There is no single true morality. There are many different moral frameworks, none of which is more correct than the others. (8)

As examples of intercultural diversity, Harman cites variations ranging from minor differences in local etiquette to differing stances on marriage and even cannibalism.

More recently, Velleman (2015) cites similar intuitions about the appearance of widespread moral diversity as motivation for his own relativist view of morality. As he puts it, the “case” for relativism consists of two “observations.” The first is that “people live and have lived by mutually incompatible moral norms” (2015: 75). The second is that “no one has ever succeeded in showing any one set of norms to be universally valid” (75).<sup>6</sup> I will not say anything further here about epistemological challenges for realism or about arguments from intercultural disagreement. I wish simply to emphasize that there are important reasons for being wary of moral realism and for embracing some relativistic alternative.

Still, there are a number of drawbacks associated with existing relativist theories. One might question whether relativism can even be formulated coherently.<sup>7</sup> I intend to press a different kind of worry. I argue next, in Section II, that well-known relativist views have difficulty accommodating, at once, a collection of three commonsense claims about the nature of moral reasons, especially as they compare to reasons of self-interested practical rationality.

It is worth asking: Might it be possible for a novel view to succeed in accommodating each of these claims, while still answering to the typical motivations for relativism? My goal is

---

<sup>6</sup> Compare also the work of Wong (2006). As support for a version of moral relativism, Wong draws attention to our experience of “moral ambivalence,” which occurs when we disagree over moral matters with people of other cultures but can appreciate the legitimacy of their alternative value systems.

<sup>7</sup> See Boghossian (2006; 2011).

to sketch an affirmative answer to that question. In Section III, I propose a novel form of normative relativism and describe its advantages. In Section IV, I elaborate upon the view by addressing a handful of especially pressing questions.

## **II. Three Theses Regarding Moral Reasons**

I take each of the following three theses about the nature of moral reasons to be highly plausible. These are claims that I suggest any meta-ethical view, realist or relativist, should seek to accommodate. These claims, briefly, are as follows. First, one cannot effectively exempt oneself from the authority of one's moral reasons by adopting contrarian attitudes. Second, moral reasons are not determined exclusively by social convention. Third, we possess both moral reasons and distinct reasons of self-interested practical rationality; these often count in favor of different behaviors.

In the rest of this section, I introduce each thesis by example and make a brief case for the conclusion that existing forms of relativism are ill equipped to accommodate all three. I focus on a handful of views that represent a diverse variety of relativist theories. I consider the forms of cultural moral relativism defended by Harman (1996) and Velleman (2015), as well as Subjectivism about normative reasons as endorsed by Street (2006) and Sobel (2016). Subjectivist views face great difficulty in accommodating thesis (1), while thesis (2) presents the greatest challenge for forms of cultural relativism.<sup>8</sup> All of these views have difficulty with (3), albeit for different reasons.

---

<sup>8</sup> I am *not* suggesting that the ability to capture the spirit of thesis (1) is precisely what renders any relativist view unable to accommodate thesis (2), or vice versa.

(1) *One cannot effectively exempt oneself from the authority of one's moral reasons by adopting unusually selfish, or merely odd, motives.*

As a simple example, suppose one has moral reason to stay home one evening to care for an ailing relative. One might develop a strong preference to go out to dinner with a friend. Intuitively, this circumstance does not shift the balance of one's reasons so that one's moral reasons are outweighed. To take a somewhat more fanciful example, one would be wrong in thinking that because one values graffiti art above all else, one therefore lacks compelling moral reason not to deface one's neighbors' property.

I am not claiming that a moral reason must always by nature trump one's reasons of other kinds. This is a stronger, and more controversial, claim.<sup>9</sup> I wish to leave open here the possibility that moral reasons may be equal in weight, or strictly incommensurable, in comparison to some other reasons. My claim is that we cannot effectively 'opt out' of moral requirements by adopting some idiosyncratic evaluative attitudes that give rise to new and opposed desire-based reasons that override the force of moral reasons. In contrast, for example, a strong preference to eat sushi for dinner *can* give rise to reasons that comparatively outweigh one's reasons for eating pizza that night. Moral reasons are relevantly different.

The task of defending this first thesis poses the most serious challenge for Subjectivist views, since according to the Subjectivist, all of a person's normative reasons arise in the same manner in connection with that person's actual attitudes. On this view, each person's normative reasons are all fixed relative to that individual's sum of attitudes. Consider Street's (2006) view, according to which the truth of any evaluative judgment of the form *X is a reason for agent A to Y* is "a function of A's evaluative attitudes," and more specifically, of whether that judgment "would be among A's evaluative judgments in reflective equilibrium" (152). On such a view, a

---

<sup>9</sup> See Wolf (1982) for one influential discussion of the relative importance of morality.

person's reason for dining out will be commensurable with a person's reason to care for a sick relative. Regardless of how it is, precisely, that the Subjectivist takes one's normative reasons to be determined by one's sum of attitudes, this approach will yield possible cases in which a person's idiosyncratic desires give rise to reasons that vastly outweigh one's reasons to behave morally.<sup>10</sup>

It will be especially difficult for the Subjectivist to accept (1) if she is committed to *Proportionalism*, the view that the force of one's reasons varies in tandem with the strength of the relevant attitude that gives rise to that reason, as well as the extent to which the behavior would promote or achieve the relevant aim or desire. On such a view, one has much stronger reason to tap dance than to keep one's promises if one has a very strong desire to tap dance but does not care much about the value of keeping promises. Sobel (2016) has argued that Proportionalism is "part and parcel" of the Subjectivist approach (306).<sup>11</sup> It might indeed prove impossible for the Subjectivist to agree with (1) without imposing *ad hoc* (and realist-sounding) constraints on the reason-giving priority of different kinds of preferences.

(2) *What one has moral reason to do is not determined solely by the standing behavioral conventions, or even laws, of one's society.*

---

<sup>10</sup> Street (2009) addresses the criticism that she is committed to the possibility of an "ideally coherent Caligula" who aims solely at maximizing others' suffering and would, accordingly, have stronger normative reasons to torture others than to behave morally. Street draws the inspiration for this character from a (1999) article by Gibbard. Street (2009) has much to say in reply; a key part of this response involves an effort to mitigate the sting of the objection that such strange agents are conceptually possible on her view. After all, Street notes, vanishingly few real people share the attitudes of this caricature. Still, the result that Caligula could have stronger reasons to torture others than to behave kindly is not a welcome one.

<sup>11</sup> Sobel's arguments are directed towards Schroeder (2007). Schroeder holds that all of a person's normative reasons are to be explained in connection with that person's desires, though Schroeder denies that the strength of a reason must be proportional to the strength of the desire and to how effectively the relevant behavior would promote the fulfillment of that desire. Sobel argues against the plausibility of this package of views.

To do what is morally right and what is conventional are not the same. One sometimes has moral reason to avoid conforming to social practices. For example, suppose the use of a hurtful derogatory term for a minority group is widespread in casual conversation within one's society. One might have moral reason to draw attention to, and to discourage, the use of that term.

The truth of (2) seems obvious. Yet certain forms of moral relativism have great difficulty accommodating it. Consider Velleman's (2015) view, according to which one's normative reasons are fixed by the way of life that is considered "ordinary" in one's community (83). According to this view, the attitudes (regarding what is good and important) and practices (like styles of dress and dancing) that we have reason to adopt are precisely those that are considered ordinary in our social group.

This is because what gives reasons their normative force on this view is the ubiquitous human social drive for mutual interpretability. Velleman emphasizes that social life, which revolves around interpersonal relationships and exchanges, is only possible for us insofar as our actions and our interests are, by and large, comprehensible to one another. This drive for mutual interpretability gives our reasons *relative* force precisely because the ways of life and ways of thinking that are considered customary vary from one culture to another.

Velleman agrees that it is possible for us to participate in social life even if our attitudes and practices are not perfectly ordinary. Yet it is important to emphasize that according to this view, we do not have *reason* to do things that are not ordinary, and so do not help to make us comprehensible to our peers. After all, the drive for mutual interpretability was proposed to be the parameter by which reasons for action derive their normative authority.

In an effort to defend the ability of Velleman's view to accommodate thesis (2), one might naturally draw attention to Velleman's claim that some communities are more "advanced" than others (2015: 97). This is so, according to Velleman, because the social drive toward mutual interpretability is a parameter in relation to which different ways of life can be judged more or less successful. However, there is an important difficulty with this suggestion. Velleman makes it clear that he takes the content of one's reasons to be fixed by the actual way of life practiced by one's community. He does not hold that the content of one's reasons is fixed by the more generally specified project of rendering oneself intelligible to one's peers. For instance, Velleman offers the example that a member of the Mbuti people has a reason for building a lean-to rather than for building a hut as a shelter. Velleman does not hold that a member of the Mbuti, or of any other community, has reason *to build the kind of shelter considered ordinary in one's community*.

There are important reasons for why Velleman might want to resist this approach. If we all, regardless of our community membership, shared the same reasons *to do the sorts of things considered ordinary in our community*, or even *to render ourselves sufficiently intelligible to our peers*, then this would no longer be a form of moral *relativism*. The latter interpretation would also yield especially counterintuitive results about the content of our moral reasons. Velleman is right not to take this approach. Yet to claim that some ways of life are more advanced than others may seem somewhat beside the point if members of a less advanced community do not possess reasons to improve their way of life.<sup>12</sup> In summary, Velleman's view provides an innovative way

---

<sup>12</sup> Velleman (2015) makes his commitment to this result clear when he states, with regard to an especially vivid example: "We Westerners are therefore in no position to say that a Kikuyu mother has reason not to circumcise her daughter – unless, that is, we can locate such reasons within the Kikuyu way of life. Even if our Western way of life is more advanced, it cannot provide reasons to the members of communities who follow different ways" (98).



to account for the relative force of moral reasons at the cost of leaving little or no room for an intuitive distinction between what is morally right and what is socially ordinary.

Harman's (1996) form of moral relativism is likely subject to similar worries, depending upon how one understands Harman's notion of the "moral framework" relative to which the truth of all moral judgments must be evaluated. On Harman's view:

For the purposes of assigning *objective* truth conditions, a judgment of the form, *it would be morally wrong of P to D*, has to be understood as elliptical for a judgment of the form, *in relation to moral framework M, it would be morally wrong of P to D*. Similarly for other moral judgments. (43, italics in original)

In other words, according to Harman, it is not strictly true that, e.g., *It is wrong for Bill to steal from his boss*. It might be true that *In relation to the moral framework of Bill's society, it is wrong for Bill to steal from his boss*. Now if a society's moral framework is simply a collection of beliefs and conventions, then Harman too will have great difficulty explaining how it would be possible for a person to have moral reason to disregard or overthrow the standing social conventions. It would be possible to insist that the moral framework of a society is something other than its actually endorsed system of moral beliefs; it might be a much-idealized version of that sum of beliefs. Recall, however, that Harman takes his view to be motivated principally by empirical considerations regarding the widespread diversity of ways of life between cultures. The less tightly connected Harman takes a moral framework to be to a description of a society's presently endorsed collection of moral beliefs, the less clear it will be that this view is supported by anthropological considerations after all. Furthermore, if Harman thinks that these two things are not very tightly linked (e.g. if one interprets the moral framework of a society as a much-idealized version of that society's actual sum of moral beliefs), then the view might then be

subject to some of the very same problems that face moral realism, such as epistemological objections. Similar worries would apply if one proposed adding (arguably *ad hoc*) supplementary restrictions to the sum of cultural beliefs that one took to give rise to true moral reasons.

*(3) We possess both moral reasons and reasons of self-interested practical rationality. These distinct kinds of reasons often call for different or even incompatible behaviors.*

I mentioned above that one might have moral reason to discourage the use of an insulting derogatory term by one's peers. Yet it seems obvious that one can also have very different kinds of reasons to do the same thing, for instance if doing so in the workplace might help to advance one's career. One might be able to garner approval from one's supervisor, for example. This is a case in which one has distinct reasons of morality and of self-interested practical rationality that count in favor of the same behavior. Alternatively, one's moral reasons and reasons of practical rationality might counsel opposing actions. For example, a farmer who lives near a lake might have practical reason to pollute the lake (because it is the most convenient way of disposing of waste from the farm), but moral reason not to do so (because the lake is used as the town's supply of drinking water).

The point that we possess both kinds of reasons, which are importantly different in character, is highly intuitive. Yet the Subjectivist and the cultural moral relativist each have difficulty in accounting for their distinctness. Velleman's moral relativism, for example, has difficulty accounting for reasons of self-interested practical rationality that come apart from moral reasons. On this view, our normative reasons derive their force in connection with the social drive toward mutual interpretability. These reasons (to act as is socially customary for us) are presumed to be moral in nature. I stressed in discussion of thesis (2) that this view will face

great difficulty explaining how we could possess specifically moral reasons to flout standing conventions. One can see, similarly, how this view will have difficulty explaining how we could have self-interested practical reasons to flout standing conventions in the service of our own desires when what we desire has, intuitively, little or no moral significance at all.

The Subjectivist, on the other hand, has difficulty distinguishing specifically moral reasons from one's normative reasons in general, all of which arise in connection with one's sum of attitudes. In discussion of thesis (1), I stressed that the Subjectivist is committed to the possibility of cases in which one's preferences give rise to reasons that vastly outweigh one's reasons to do what is intuitively the morally right thing. Another related worry for the Subjectivist is that certain possible agents would lack moral reasons entirely, insofar as they failed to hold relevant attitudes that would give rise to such reasons in the first place. In anticipation of this and similar objections, Street notes that her view makes room for the possibility that all psychologically normal human beings will possess reasons to treat others decently.<sup>13</sup> Still, however, it is worth trying to resist the result that people would not possess any moral reasons if they lacked the relevant attitudes.

### III. Group Agency Meets Meta-Ethics

In this section, I outline a novel account of moral reasons as well as reasons of self-interested practical rationality. I argue this view can make progress in accommodating all three claims from Section II while still holding out the promise of delivering on the initial motivations for relativism mentioned in Section I. First, a caveat. It is by no means my ambition to convince

---

<sup>13</sup> See Street (2012), especially Section 7, "Morality and contingency," and Street (2016), especially Section 12.11. Schroeder (2007) makes a related point in defense of his view that all of one's normative reasons are to be explained in connection with one's desires. On his view, there may be universal reasons if there are relevant desires that all human beings possess. See Schroeder's discussion of "too few reasons" in Chapter 6 of *Slaves of the Passions* (2007).

the reader that my proposal is the only viable meta-ethical option. Instead, it is my aim to investigate, by example, how we might make progress in answering some of the most difficult objections to relativism by revisiting our assumptions about how a relativist view must look. I would be equally delighted if my work here inspired the reader to pursue a different approach as a way of yielding new answers to old questions about the limitations of normative relativism.

According to the proposal, as well as for existing forms of relativism, all of the normative facts are attitude-dependent. This claim is important, since this tight link between the normative facts and contingently held attitudes is the key feature of relativist views that enables them to sidestep epistemological objections to moral realism and to account for the possibility of cultural differences in moral norms, as discussed in Section I. I take all normative facts to be grounded in (but not logically equivalent or reducible to) contingent facts about actually existing agents, including facts about their aims (including goals and values as well as desires) and about the capacities of those agents as well as the limitations imposed by their practical circumstances. What makes the view unique is its insistence that there are a wider variety of agents, beyond just persons, whose attitudes give rise to attitude-dependent reasons.

Persons are, of course, one sort of agent. I hold that reasons of self-interested practical rationality apply directly to individual persons in light of their aims. Yet on my view, moral reasons are importantly different in kind. I hold that what distinguishes moral reasons is that they apply directly to whole societies understood as agents in their own right (and derivatively to individual persons). The content of moral reasons is determined by aims attributable to whole societies.

### ***III.1 Expanding our Understanding of Agency***

There is precedent, and independent motivation, for thinking that plenty of the kinds of attitudes (such as beliefs, goals, aims, and desires) that are most often attributed to persons, and taken to provide them with reasons, are equally attributable to groups understood as agents in their own right. List and Pettit (2011) argue that groups such as corporations, states, and organizations, count as genuine agents. They can act, on the basis of information available to them, in pursuit of their interests. Recognizing a group as an agent need not involve attributing to it a unified, conscious mind. What agency does require, on List & Pettit's account, are the following three features: "representational states" that represent features of the agent's environment, "motivational states" that specify how the agent requires for things to be in its environment, and the possession of the capacity to process those representational and motivational states in such a way as to "intervene suitably" in its environment as needed (20). These representational and motivational states attributable to group agents play functional roles analogous to the roles of the mental states of beliefs and desires in human beings.

Gilbert (2013) also defends the idea that some of the attitudes, such as beliefs, that we most often attribute to persons can also be attributed directly to groups of persons acting together. Gilbert argues that it is neither necessary nor sufficient for the majority of the members of a group to possess individual tokens of a belief in order for the group to hold that belief collectively (Gilbert, 2013: 168). Gilbert offers the example of a legal opinion officially held by a court. The opinion, in its final form as held by the whole court, is not an opinion held in the same form by any of its individual justices (2013: 170).

As part of my own proposal, I endorse a functionalist characterization of the requirements of agency, which draws greatly from List & Pettit's view, but with a crucial modification. I hold that an agent not only possesses aims, or "motivational states," but has the more demanding capacity (also understood in functional terms) to *self-determine* its own aims. I deny that things like coffee makers or heat-seeking missiles are genuine agents, even though such objects arguably fulfill List & Pettit's characterization of the requirements of agency, including the possession of motivational states that guide their behavior.<sup>14</sup> Yet they do not select their own aims. Correspondingly, I hold they are not agents and do not possess normative reasons. I hold that the following three conditions are each necessary, and jointly sufficient, for status as an agent. An agent has the abilities to (i) determine its aims, (ii) represent information about itself and/or its environment, and (iii) act in pursuit of its selected aims by modifying itself and/or its environment.<sup>15</sup>

It is easy to see how these points about agency might extend to societies. A society can form aims, like the aim of protecting the health of its children or of preserving its cultural heritage. Societies can also hold beliefs; for example, it might become known by a society that lead paint harms developing children, or that a holiday or festival occurs at a certain time each year.<sup>16</sup> A society can also engage in actions in pursuit of its aims; for example, a society seeking to protect its heritage might promote the teaching of certain songs and stories. Or a society might discourage parents from leaving young children unattended through patterns of social sanctions.

---

<sup>14</sup> The example of a heat-seeking missile is from O'Brien (2015).

<sup>15</sup> My conditions (ii)-(iii) involve another important departure from List & Pettit (2011). They require an agent to have "the capacity to process its representational and motivational states, leading it to intervene suitably in the environment whenever that environment fails to match a motivating specification" (20). Yet I hold that the ability to modify oneself as well as one's environment is a hallmark of agency. For instance, people may seek to acquire new skills if necessary.

<sup>16</sup> If the attribution of "beliefs" to groups sounds implausible, I am just as happy to say that groups have representational states that play the same functional role as beliefs; we can reserve the term "belief" for creatures with unified, phenomenally conscious minds.

### **III.2 The Proposal in More Detail**

I propose that the different kinds of normative reasons that we recognize intuitively, including moral reasons and reasons of self-interested practical rationality, are distinguished by the fact that they apply directly to different sorts of agents. (We might say that kinds of normative reasons are *relative* to different sorts of agents.) For each kind of reason, the content of those reasons – what it is in particular that they call upon the agent to *do* – is fixed in connection with the aims of the relevant agent. I take moral reasons to be the kind of reason that applies directly to societies in light of their aims. Diverse societies may therefore possess different moral reasons in light of their different aims. Moral reasons are distinct from self-interested reasons of practical rationality, which apply directly to persons in light of their own individual interests. These reasons vary from person to person.

Despite what the terminology of “self-interested” suggests, one’s reasons of practical rationality need not be purely selfish. One’s personal aims may include, e.g., pro-social and altruistic values that give rise to practical reasons to uphold those values.<sup>17</sup>

I propose that normative claims are true when they correctly describe a relationship between an agent’s aim(s) and the means available to the agent for pursuing those aim(s), given the capacities of that agent and its practical limitations in context.<sup>18</sup> Consider the following example:

- (a) Mary has a reason<sub>practical rationality for Mary</sub> to notify the post office of her change of address.

---

<sup>17</sup> In what follows, I will sometimes speak of “practical” reasons for short. I introduce the longer term to distinguish these reasons from moral ones; other authors might use the term “practical” to designate all reasons for action.

<sup>18</sup> I suspect that a similar treatment could be given of normative “requirements,” “prohibitions,” or “oughts,” if preferred, so long as one took a roughly instrumental view of these other normative notions. It is not my intention to suggest that reasons must be normatively fundamental. See Wodak (this volume).

Claim (a) will be true so long as (i) the action of notifying the post office of her change of address is an action of the kind that Mary has the capacity to perform in this context, and (ii) that action is something that would facilitate the pursuit of an aim that Mary holds. The most obvious candidate here for a relevant aim, which would seem to be implicit in any statement of this claim in ordinary conversation, is Mary's goal of receiving her mail. Since these notions of (i) having a capacity, and (ii) facilitating an aim, admit of degrees, we can also make sense of the commonsense idea that reasons vary in force or strength. For example, we can say that Mary has stronger reasons to take actions that are more efficient means to her aims. So Mary might have stronger reason<sub>practical rationality</sub> to complete a change of address form online tonight rather than in person on Monday, assuming that it is easier for Mary to submit a form online than to walk to the post office, or that getting the news to the post office sooner would help Mary to recuperate more of her mail. In addition, the force of her reasons may vary in strength in connection with the importance that she attributes to the relevant aim.

Analogous points will hold for the aims and reasons of societies. I have offered the examples of societies aiming to protect the health of their children or to preserve their cultural heritage. Again, there will be better and worse ways, providing the basis for stronger and weaker reasons for action, of pursuing these aims.

I have mentioned that according to this proposal, reasons of morality apply directly to societies, rather than persons. I take this to be the right result, for two reasons. First, facts about moral reasons will be grounded in, among other facts, facts about the attitudes attributable to societies as a whole, which its individual members may or may not share. Second, the moral reasons that apply directly to societies will often be reasons for actions of the type that only a whole society, and not an individual person, could feasibly carry out. Consider the example of



holding an annual festival; this is not an action that one person could perform. Thus it hardly makes sense to say that people and societies possess the very same reasons for action.

Yet I hold that there is a sense in which moral reasons apply *derivatively* to the human members that constitute the society to which those reasons apply directly. To be precise, according to the proposal:

A member, *m*, of society *S* will have derivative moral reason to play *m*'s part in *S*'s effort to  $\phi$  if and only if *S* has moral reason to  $\phi$ .

Thus for *m* to possess a derivative moral reason requires that *m* is a member of *S* and that *S* possesses some moral reason in light of *S*'s aims. This does not require that *m* has any particular attitudes whatsoever. For example, *m* need not endorse either *S*'s aims or the state of *m*'s own membership in *S*.

What it is for a member, *m*, to play his or her "part" in *S*'s effort to  $\phi$  will depend upon *m*'s role within *S*. For example, it will be appropriate for a political leader and a teenaged student to act very differently in accordance with a society's effort to preserve its cultural heritage, or to keep its members safe. If a society has reason to prepare for an impending natural disaster, this will result in, e.g. a mayor, an ambulance driver, and an elderly citizen having derivative reasons to engage in very different behaviors as part of the society's effort. I describe derivative reasons as counting in favor of *m* playing *m*'s "part" in the society's effort precisely because action at the level of the group agent involves a number of members contributing to that action. Note that by "S's effort to  $\phi$ " I mean to indicate the group action by *S* as a whole that the group has reason to undertake in order to effectively pursue its aim(s). Thus I do *not* hold that *m*'s playing *m*'s part involves doing whatever *m* can do individually to maximize the chance of *S*'s achieving its aims,

given the actual circumstances of whether or not other members are playing their part. To hold this would shift unreasonable moral burdens onto individuals.

### ***III.3 The Payoff: Revisiting Claims (1)-(3)***

In this section, I clarify how this proposed view of moral and practical reasons makes it possible to accommodate all three theses from Section II. Consider the first thesis:

*(1) One cannot effectively exempt oneself from the authority of one's moral reasons by adopting unusually selfish, or merely odd, motives.*

According to the proposal, people possess derivative moral reasons whenever they are members of a society that possesses aims, giving rise to moral reasons that apply directly to the society as a whole. Suppose a society has the aim of protecting the health and safety of its children. This will give members of the society derivative moral reasons to do different things, depending on their roles within that society. For example, parents will have derivative moral reasons to ensure that their own children have access to food, warm clothing and medical care. Others will, less demandingly, have derivative moral reasons to do things like pay enough in taxes to support basic emergency services for children. They will also have derivative moral reasons to refrain from doing things that directly harm children.

Since moral reasons arise in connection with the aims held by a society as a whole, derivative moral reasons will apply to all members of a society regardless of the idiosyncratic aims of any particular member, no matter how strongly that individual's desires or preferences are held.<sup>19</sup> It is never true, on my view, that one's reasons of self-interested practical rationality,

---

<sup>19</sup> It is not strictly true that one's idiosyncratic psychology makes zero difference in fixing the moral reasons that apply to one's society, since on this view, facts about the society as an agent supervene upon facts about all of its members. In the case of a whole society, which I am discussing here, one person's idiosyncratic preferences will make virtually no difference in determining the content of the society's moral reasons, which will suffice to defend

which are determined by one's own interests, can effectively outweigh one's derivative moral reasons. (In order to exempt oneself from the force of one's derivative moral reasons, one would need to cease to be a member of the society altogether.) On my view, no such comparisons of the weights of different kinds of reasons are strictly possible, since different kinds of reasons apply directly to different sorts of agents, and no agent is objectively privileged.

Next, skipping ahead, we can easily see how the proposed theory will be able to accommodate the third thesis:

*(3) We possess both moral reasons and reasons of self-interested practical rationality.*

*These distinct kinds of reasons often call for different or even incompatible behaviors.*

On the proposed view, moral and practical reasons are distinct in kind. Moral reasons are the reasons that apply directly to societies, and derivatively to their human members, in light of the societies' aims, while practical reasons are the reasons that apply directly to individual persons in light of their own aims. The view vindicates the idea that human beings simultaneously possess reasons of both kinds. A person will possess practical reasons in light of any aims she adopts. A person will also possess derivative moral reasons as long as she is a member of a society.<sup>20</sup> A person need not possess any particular attitude in order to be subject to derivative moral reasons. Unsurprisingly, moral and practical reasons may count in favor of different or even incompatible actions, since they arise in light of different aims held by distinct agents, and so may call for different actions as means to those aims.<sup>21</sup>

---

thesis (1). Yet in a case in which a group that is not a typical society were composed of very few members, say two or three, then one person's idiosyncratic attitudes would make a significant difference regarding the reasons possessed directly by that group. Thank you to Selim Berker for pressing this point.

<sup>20</sup> Note that a person may be a member of multiple societies, which may or may not overlap. A member of multiple societies will possess distinct moral reasons of multiple kinds.

<sup>21</sup> In some cases, a person's practical reasons and her derivative moral reasons will correspond, giving her distinct reasons of two different kinds for behaving in the same manner. (This does not require that the person and her society each endorse the same aim.) Of course, an individual might personally endorse some or all of her society's aims; this is one case in which her practical and derivative moral reasons would correspond.

The proposal's appeal to the idea of group agency (and more specifically to the claim that societies are distinct agents from each of the people that constitute them) is thus what makes it possible to insist upon a clear and principled distinction between the nature and source of moral versus practical reasons, which some of the other relativist views I have discussed struggle to provide. The claim that each of us is an individual agent as well as part of society is, furthermore, what makes it possible on this view to insist that we are simultaneously subject to both moral and self-interested practical norms.

Finally, consider the second thesis:

*(2) What one has moral reason to do is not determined solely by the standing behavioral conventions, or even laws, of one's society.*

The proposed view delivers the desired result that the standing conventions in a society do not necessarily reflect what is morally appropriate, though they might, in whole or in part. According to this view, moral facts are grounded in facts about the aims and capacities of societies as well as about their practical circumstances. It should come as no surprise that a real society might be incredibly inefficient or downright self-undermining in pursuit of its aims; for example, a society might endorse political principles upholding the equality of all citizens while the social norms reflect a deeply entrenched class system or bigotry. A society might also hold a variety of aims that are at odds. The proposal can make sense of the idea that a society may have moral reasons, and even decisive moral reasons, for overthrowing social conventions in order to more efficiently pursue the values that the society really holds most dear. The view does not imply that moral reasons must be reasons to uphold the standing conventions.

One might worry that a close cousin of the original problem for relativism arises for the proposed view. The view may not be committed to the claim that whatever is customary is

morally right; yet the proposal will be committed to the claim that what is morally right within a society is ultimately fixed by the aims of the society. This is the first of four objections that I consider in the next section.

#### IV. Selected Questions & Objections

1. *What if a society has terrible aims? Do its members really have derivative moral reasons to assist in achieving those aims?*

It follows from my view that a society possesses moral reasons to take effective means to the aims that it possesses as a contingent matter of fact. Before I say more, I find it important to emphasize a point similar to one made by Sobel (2016: 28) in defense of the Subjectivist. I am not, by recognizing that a society could have genuine reasons to engage in certain actions, asking anyone to withdraw their reaction of condemnation. We might indeed have practical reasons (or our society might have moral reasons, or both) to interfere and attempt to stop other agents from doing what they have genuine reason to do, for example if an agent has reason in light of its aims to engage in violence directed towards us. We might also have practical reasons to join or leave various groups.

I also find it important to draw attention to the relativistic features of the view. To claim that a society has moral reason to  $\phi$  is by no means, on my view, an objective vindication of the appropriateness of  $\phi$ -ing. On my view, there are no such facts (from an ultimate or “God’s eye” perspective) about what all possible agents must have all-things-considered reason to do. Instead, there are only facts about what particular agents have reason to do in light of their contingently held aims.

Furthermore, my view has some special features that other relativist views lack that can help to get more satisfying results about moral reasons in cases in which a society has seemingly atrocious aims. I think it pays to be careful about the questions of what the aims of the society in question really are, and about what are the most effective means to those aims. In order to illustrate this point, I will offer an example that is admittedly favorable to my view. We know that certain historical communities, such as the Aztec, Maya, and Inca civilizations, engaged in the ritual sacrifice of living people, including children.<sup>22</sup> This, of course, strikes us as atrocious. Now suppose the Mayans took themselves to be doing this in order to appease certain gods or spirits, and to thereby accrue certain benefits, such as favorable weather or farming conditions. The Mayans may really have aimed at appeasing the gods. Yet if no such gods existed (and had no control over the weather), then according to my view, the Mayan society did not have moral reason to sacrifice anyone, since such an action would not have been an effective means to that aim. This yields much more plausible results than the view that the Mayan's moral reasons were fixed simply by custom or convention, for example.

I also wish to emphasize that the true aims of a society cannot always be “read off” of explicit statements of social aims, such as a leader's decrees, a written constitution, or even a majority vote.<sup>23</sup> To identify the true aims of one person may be similarly difficult; a person might lie, or simply be confused, about her intentions. Any attempt to identify the aims of any

---

<sup>22</sup> This fictionalized example is inspired by a 2018 *National Geographic* article by Kristin Romey describing an archeological discovery of the remains of human sacrifice.

<sup>23</sup> Fact about the properties of societies, including their aims, will supervene on facts about its human members. Yet I deny that the society's aims must be determined by some aggregation function of the aims of its members. (Majority vote on an issue provides one simple example of an aggregation function.) Compare List & Pettit (2011), who discuss a number of problems for various aggregation functions, though they are sympathetic overall to the idea that group attitudes are determined by aggregation of member attitudes. Compare also Copp (2001). Copp holds that a moral code is justified for a society just in case the society would be rationally required to select that code in preference to any alternative (2001: 104). For Copp, a society prefers a moral code just in case that code is an option for the society and the members who have an attitude toward that option are “nearly unanimously in favor” of it (2001: 148). Note also that for Copp, a person may not necessarily have any reason to subscribe to a moral code even if that moral code is justified relative to that person's society (2001: 243). I hold that people possess derivative moral reasons whenever they are members of a society, regardless of their own aims.

agent, as part of a functionalist approach, will require a more complex process of observing the agent's behavior over time in combination with a theoretical project of modeling or reconstructing the agent's sum of intentional states in an attempt to explain that behavior.

One might worry that the objection I face regarding a society's objectionable aims is largely the same as the one that arises for the Subjectivist, viz., that a person may lack reasons to behave morally if she fails to hold the relevant attitudes. Yet I think important progress has been made. On the view I propose, we have the resources to distinguish between practical and distinctively moral reasons, and to insist that all members of a society will possess moral reasons regardless of their idiosyncratic attitudes. The Subjectivist has difficulty in carrying out both of these tasks. This first objection remains one of the greatest challenges for my view. I hope I have shown that the view I propose has new resources to make progress in dealing with this sort of objection in comparison to the Subjectivist and cultural moral relativist.

*2. I grant that societies may be agents. Why, then, would a society not possess both moral reasons and reasons of practical rationality, like all agents?*

I use the term "moral" within my framework to refer to the kind of reasons that apply directly to societies in virtue of their aims. I use the term "practical" for the reasons that apply directly to persons in light of their aims. There is a sense in which moral reasons simply are practical reasons for societies; societies possess moral reasons to take what are, in context, instrumentally efficient means to the aims they hold. Thus on my view, societies do possess reasons to effectively pursue their interests, which one might otherwise have been inclined to call reasons of practical rationality.

I also hold that whole societies, just like persons, might also possess derivative moral reasons that they inherit from even larger groups of which they are members. There may be cases in which a small society (e.g. an ancient Greek city-state, or a rural American town) is a member of a larger society (corresponding roughly to Greece, or West Virginia). In such a case, there will be multiple distinct kinds of moral reasons - those that apply directly to the larger society and are fixed by its interests, and those that apply directly to the smaller in light of its own aims. Both would apply derivatively to individual human members. For example, West Virginia society as a whole could possess moral reasons to distribute rights to natural resources, like water or timber, equitably between towns.

There must be some societies that are not members of any additional, larger society that qualifies as a group agent. (I do *not* claim that all people on Earth are united as a group agent.) These societies will not inherit any additional derivative moral reasons. I take this to be a welcome result, since moral reasons are, very roughly speaking, reasons for a group to promote a shared and stable way of life for the mutual benefit of its members. If a whole society is not itself participating in a way of life shared with other coordinating group agents, then that society will not inherit derivative moral reasons. It will possess only the moral reasons that apply to it directly in light of its interests and therefore derivatively to each of its members, who are engaged in a shared way of life.

*3. If moral reasons apply directly to societies in light of their aims, then why do we as individuals also possess “derivative” moral reasons?*

I hold that the reasons that apply directly to societies also apply derivatively to the human members of the group for the simple reason that a group agent is composed of human members. I



do not claim that members inherit derivative reasons in light of their endorsement of the group's aims, or in virtue of choosing to endorse a relationship with the group agent. Membership is not a contract. Instead, a member is literally a part of a group agent. I do not conceive of the membership relation between a society and a member as a relation between two full and separate agents at all. The way in which a person possesses the group's reasons is subject to some qualifications (hence the notion of a "derivative" reason) only because of the point that one person is a very small part of the group agent as a whole.

It is a consequence of taking the idea of group agency seriously that human beings often form a part of many different agents at once. If groups can be genuine agents in their own right, then agents are not discrete and separable. It is not true that each agent corresponds to a single body or a single conscious point of view. Agents overlap. Each person, as one agent, possesses her own practical reasons. Yet a human being, *qua* group member, also possesses derivatively the reasons of the group agent as a whole.

While this might sound unusual, I think it captures the commonsense intuition that there are genuine normative dilemmas between what is best for us as individuals (reflected here by practical reasons) and best for society (as reflected by moral reasons). The aims of these different agents are different sources of reasons.

*4. In support of the idea that societies can be agents, you cite literature that defends the claim that a whole range of groups can be agents. Do you think other groups, like corporations and philosophy departments, are agents, too? Do they possess moral reasons?*

I embrace the result that these groups could possess reasons of their own, so long as such groups met my requirements for agency. The reasons that apply directly to corporations, or to

clubs, would not be moral reasons on my view, however. Given my strategy of associating different kinds of normative reasons with different sorts of group agents, I take moral reasons to be the kinds of reasons that apply to societies, as opposed to some other group, because I suspect that this view best accommodates our intuitions about the content of moral reasons. That is, our moral reasons are largely other-regarding; our moral reasons tell us how to govern our cooperative behaviors and relationships with others in persisting social groups, including our family members, our friends, and other members of our community. I take it that the kinds of aims that are most plausibly attributed to societies, of any existing group agent, are most appropriate as the basis for moral reasons. The reasons possessed by a group such as a sports team in light of its aims, and possessed derivatively by its members, would be an entirely different kind of reasons.

Note that not all “social groups” will qualify as group agents as I defined them. For instance, Ritchie (2018) defends a social ontology that makes it possible to recognize categories like “Asian women” and “middle-class men” as “social groups.” If these groups do not self-determine aims and pursue them as unified entities, then on my view they are not agents and there will be no distinct kind of normative reason that applies exclusively to e.g. Asian women or middle-class men.

## **V. Concluding Remarks**

The view I defend here combines an endorsement of the relativist thesis that all normative facts are attitude-dependent with a more permissive view of what can count as an agent. By recognizing that societies as well as persons can be agents with aims of their own, and by associating moral and practical reasons with the different sorts of agents exemplified by

societies and persons, this view seeks to offer a new approach to normative relativism that more faithfully captures the distinctive natures of moral and practical reasons. I hope the reader will find that the view sketched here has some novel resources that can make progress in avoiding some of the features that make existing forms of normative relativism unsatisfying, while still holding out the promise of delivering on the initial motivations for such views.<sup>24</sup>

---

<sup>24</sup> For helpful discussion, I wish to thank audiences at the 2018 Madison Metaethics Workshop, the 2018 CEPPA Graduate Conference at the University of St. Andrews, and New York University's Philosophy Department. Thanks especially to Paul Boghossian, Sharon Street and Paul Horwich.

## References

- Ayer, A.J. 1936. *Language, Truth, and Logic*. London: Gollancz.
- Bedke, Matthew. 2009. "Intuitive Non-Naturalism Meets Cosmic Coincidence," *Pacific Philosophical Quarterly*, 90(2): 188-209.
- Bedke, Matthew. 2014. "No Coincidence?\*" in Russ Shafer-Landau (ed.), *Oxford Studies in Metaethics*, Vol. 9, 102-125. Oxford: Oxford University Press.
- Berker, Selim. 2014. "Does Evolutionary Psychology Show That Normativity is Mind-Dependent?" in Justin D'Arms and Daniel Jacobson (eds.), *Moral Psychology and Human Agency: Philosophical Essays on the Science of Ethics*, 215-252. Oxford: Oxford University Press.
- Blackburn, Simon. 1993. *Essays in Quasi-Realism*. New York, NY: Oxford University Press.
- Blackburn, Simon. 2005. "Antirealist Expressivism and Quasi-Realism," in David Copp (ed.), *The Oxford Handbook of Ethical Theory*, 146-162. Oxford: Oxford University Press.
- Boghossian, Paul. 2006. "What is Relativism?" in Patrick Greenough and Michael P. Lynch (eds.), *Truth and Relativism*, 13-27. Oxford: Clarendon Press.
- Boghossian, Paul. 2011. "Three Kinds of Relativism," in Steven D. Hales (ed.), *A Companion to Relativism*, 53-69. West Sussex: Blackwell Publishing.
- Copp, David. 2001. *Morality, Normativity, and Society*. New York, NY: Oxford University Press.
- Copp, David. 2008. "Darwinian Skepticism about Moral Realism," *Philosophical Issues*, 18, *Interdisciplinary Core Philosophy*, 186-206.
- Dyke, Michelle M. 2019. "Bad bootstrapping: the problem with third-factor replies to the

- Darwinian Dilemma for moral realism,” *Philosophical Studies*, forthcoming, Online First  
08 May 2019. <https://doi.org/10.1007/s11098-019-01301-4>
- Enoch, David. 2010. “The epistemological challenge to metanormative realism: how best to understand it, and how to cope with it,” *Philosophical Studies*, 148, 413-438.
- Enoch, David. 2011. *Taking Morality Seriously: A Defense of Robust Realism*. Oxford: Oxford University Press.
- Gibbard, Allan. 1999. “Morality as Consistency in Living: Korsgaard’s Kantian Lectures,” *Ethics*, 110, 140-164.
- Gilbert, Margaret. 2013. *Joint Commitment: How We Make the Social World*. Oxford: Oxford University Press.
- Harman, Gilbert, and Judith Jarvis Thomson. 1996. *Moral Relativism and Moral Objectivity*. Cambridge, MA: Blackwell.
- Joyce, Richard. 2006. *The Evolution of Morality*. Cambridge, MA: The MIT Press.
- Joyce, Richard. 2013. “Irrealism and the Genealogy of Morals,” *Ratio*, 26(4): 351-372.
- Joyce, Richard. 2016. “Reply: Confessions of a Modest Debunker,” in Uri D. Leibowitz and Neil Sinclair (eds.), *Explanation in Ethics and Mathematics*, 124-145. Oxford: Oxford University Press.
- Kitcher, Philip. 2005. “Biology and Ethics,” in David Copp (ed.), *The Oxford Handbook of Ethical Theory*, 163-185. Oxford: Oxford University Press.
- List, Christian, and Philip Pettit. 2011. *Group Agency: The Possibility, Design, and Status of Corporate Agents*. Oxford: Oxford University Press.
- Mackie, J.L. 1977. *Ethics: Inventing Right and Wrong*. New York, NY: Penguin Books.
- O’Brien, Lilian. 2015. *Philosophy of Action*. Houndmills, Basingstoke, Hampshire: Palgrave

Macmillan.

Ritchie, Katherine. 2018. "Social Structures and the Ontology of Social Groups," *Philosophy and Phenomenological Research*, Early View 20 September 2018.

<https://doi.org/10.1111/phpr.12555>

Romey, Kristin. 26 April 2018. "Exclusive: Ancient Mass Child Sacrifice May Be World's Largest," *National Geographic*. <https://news.nationalgeographic.com/2018/04/mass-child-human-animal-sacrifice-peru-chimu-science/>

Schechter, Joshua. 2010. "The Reliability Challenge and the Epistemology of Logic," *Philosophical Perspectives*, 24, *Epistemology*, 437-464.

Schroeder, Mark. 2007. *Slaves of the Passions*. New York, NY: Oxford University Press.

Setiya, Kieran. 2012. *Knowing Right From Wrong*. Oxford: Oxford University Press.

Shafer-Landau, Russ. 2003. *Moral Realism: A Defence*. Oxford: Oxford University Press.

Shafer-Landau, Russ. 2012. "Evolutionary Debunking, Moral Realism and Moral Knowledge," *Journal of Ethics & Social Philosophy*, 7(1): 1-37.

Sobel, David. 2016. *From Valuing to Value: Towards a Defense of Subjectivism*. Oxford: Oxford University Press.

Street, Sharon. 2006. "A Darwinian Dilemma for Realist Theories of Value," *Philosophical Studies*, 127, 109-166.

Street, Sharon. 2009. "In Defense of Future Tuesday Indifference: Ideally Coherent Eccentrics and the Contingency of What Matters," *Philosophical Issues*, 19, *Metaethics*, 273-298.

Street, Sharon. 2012. "Coming to Terms with Contingency: Humean Constructivism about Practical Reason," in James Lenman and Yonatan Shemmer (eds.), *Constructivism in Practical Philosophy*, 40-59. Oxford: Oxford University Press.

- Street, Sharon. 2016. "Objectivity and Truth: You'd Better Rethink It," in Russ Shafer-Landau (ed.), *Oxford Studies in Metaethics*, Vol. 11, 293-334. Oxford: Oxford University Press.
- Velleman, J. David. 2015. *Foundations for Moral Relativism: Second Expanded Edition*. Cambridge, UK: OpenBook Publishers. <https://dx.doi.org/10.11647/OBP.0086>
- Wolf, Susan. 1982. "Moral Saints," *The Journal of Philosophy*, 79(8): 419-439.
- Wong, David. 2006. *Natural Moralities: A Defense of Pluralistic Relativism*. Oxford: Oxford University Press.