

On Being Inside Social Morality and Seeing It

Gerald Gaus

Eric Mack's "Inside Public Reason" is thorough and fair-minded review of *The Order of Public Reason*. My deep thanks to him for his insights, as well as his judiciousness. In these remarks I cannot take up all the important matters he raises; in particular I put aside two important issues — the analysis of the political and discussion of how contingent social processes play a fundamental role in public justification.¹ I plan to take up the latter on another occasion.

1. ON NOT SEEING WHAT IS IN FRONT OF US

As Mack observes, *The Order of Public Reason* is a long book. It is long partly because, in comparison to most works in social and political philosophy, it employs a rather wide array of tools and methods to model what a free social morality would look like. However, it is also long because in an odd way its subject matter became invisible in late twentieth/early twenty-first moral philosophy, and much of the book is devoted to trying to get us to see it again. I say "again" because earlier moral philosophers such as the Hume, Hegel, Mill,² and more recently Kurt Baier and P.F. Strawson, saw it clearly.³ And so did John Rawls.⁴ But though Rawls focused on it, most of his interpreters and students did not; it remained, if not totally invisible, shrouded, to them.

Social morality is so hard for philosophers to see as it occupies a place between the moral judgment of the individual autonomous agent and the political. In much political philosophy these exhaust the space of normative thinking. For many all moral thinking is about searching for, and acting on one's best judgment of, moral

truth. Moral philosophy thus understood is an inquiry into right and wrong, good and bad, and so on. An individual who seeks to act morally is one who comes to her soundest judgment about what is right and wrong, good or bad, and acts in accordance with it. For some this is based on the epistemic task of truth discovery, for others it is about a certain mode of universalizable deliberation. The common supposition is that to be a moral agent is to act on one's judgment of what is morally true, or morally best.

Many believe that the judgment of the individual exhausts the sphere of the moral, and so advocate a sort of "philosophical anarchism" or extreme libertarianism, insisting that moral agents must act on morality as they understand it, and only some sort of explicit authorization by a moral agent could justify alienation of this moral sovereignty to the political. Thus the great suspicion of political authority in so much contemporary moral thinking. After all, if an individual's judgment is about what morality requires of her, what room is there for any other requirement that might conflict with this? Surely morality is the supreme leader of us all. Even those who seek to legitimate political authority tend to place the center of gravity of moral thinking on moral truth. Genuine political authority, some maintain, must act in the service of moral truth.⁵ Others make the recognition of such authority itself part of the moral truth.⁶

The social contract theorists saw the political as a way to adjudicate the moral. Hobbes, Locke, Rousseau and Kant brought to the fore a crucial problem for social life among autonomous moral agents under modern conditions: using their reason as well as could reasonably be expected, people will arrive at conflicting judgments about morality.⁷ The solution of the social contract theorists was public reason as expressed by the political umpire.⁸ When using our private reason we disagree about what is right and good, and so will come into conflict; we thus erect political

institutions that articulate public reason, providing public guidelines for coordinated action, and public adjudication of disputes.

The Order of Public Reason departs from this view of the terrain of the normative in two ways. First, and most fundamentally, it tries to bring back into our vision the realm of social morality: the shared rules, shared expectations and requirements, and shared understanding of appropriate demands and sanctions that structure our life together, and which are neither simply inferences from individual moral judgment nor the dictates of the law. Recent work on social norms and non-state institutions are bringing this structure into focus as the subject of investigation,⁹ though philosophers typically resist the claim that this is a moral realm distinct from the individual and the political. *The Order of Public Reason* seeks to show that it is. The book's second departure is to show how this realm of social morality can be reconciled with individual moral judgment and its claim to sovereignty over one's own actions. *The Order of Public Reason* does not seek to arrive at yet another theory of moral truth and correct moral individual judgment; it seeks to show how a social moral framework can arise among those committed to their diverse theories of moral truth and rightness. While no one's vision of the moral truth can dictate the social moral framework, the aim is for all to see the social moral framework as normative, given their individual understandings of moral truth.

2. OF BILLY CLUBS AND AUTHORITY

Because *The Order of Public Reason* seeks to focus on this often unseen realm of social morality, and because it is reasonable to suppose that the book must be analyzing something familiar, it is tempting to interpret it as if it is yet another treatise about

moral rightness and wrongness *per se*. I believe that Mack, despite his care, sometimes does so. He thus writes:

One needs to emphasize that Gaus's essential concern is with justifying the authority of *morality* and reconciling the authority of *moral* rules with autonomy. Gaus is impressed both by the power of moral prescriptions to guide our conduct and the need for our conduct to be guided by such rules if we are to achieve mutually beneficial coordination. *Indeed, with respect to the power of moral prescriptions, Gaus seems to hold that one's deployment of such prescriptions to get others to act as one desires is at base as problematic and in need of justification as is one's deployment of billy clubs.* (Mack ms. p. 2, emphasis added)

Remember the fundamental claim of social contract theory: in a diverse world we employ our reason to arrive at different moral judgments. Now the question arises: what happens when I engage in a social life with others who do not concur with my moral judgments? I conclude they must not use their billy club on me, but they see nothing wrong with doing so. Leaving aside for now the appeal to the political, I can simply insist on my own, controversial, conclusions and insist that they must do as I tell them, and perhaps seek to enforce my claim — I say that I am morally justified in using a billy club against them. I assert that my view of justice should rule them. Thus we have Kant's state of nature, in which each claims "the right to do what *seems just and good to him*, entirely independently of the opinion of others."¹⁰

From both an objective and a subjective point of view it would be a great boon to share a social normative framework — one in which we concur on the claims we are justified in making against each other. From the objective point of view, cooperation among strangers in a diverse society is immeasurably enhanced when, confronting another simply as a stranger, we share normative expectations, and understanding

of what can be demanded and the appropriate response to noncompliance. In lieu of such a framework, we are unlikely to coordinate our moral expectations. Since we do not know whether they share our understanding of morality (they probably do not, if this means an outlook on what is morally best or most just), we do not know when we will have uptake of our demands, and when we can expect resistance. Mack is a Lockean-inspired libertarian; when he confronts the stranger in the street, he cannot suppose that the stranger sees things in the same way. We all need to have a good grasp on the normative expectations of others, and have confidence that they will conform to expected courses of action. Without a social morality, we are thrown back on the political as our only way to adjudicate moral demands and set expectations.

Subjectively, it is deeply different to issue a moral prescription that we believe to be right but which we know others do not accept and can see no reasons with which to comply, and a prescription that we understand to be grounded in shared rules that both we and others endorse as normatively justified. In the former case, I insist that I have superior insight to others, and they should be guided by my reason, though I realize that from their perspective they will see this as an untoward demand. When we all embrace a common rule as normative, as I argue in *The Order of Public Reason*, the reactive attitudes of indignation and resentment are well-grounded. Because we share normative expectations that each sees as justified, we share a rich moral life with far-flung others. The difference is between relying on my self-appointed authority to tell others what to do, and drawing on the shared authority of publicly endorsed rules.

So our question is: what rules governing billy clubs can we all recognize as genuinely normative? *The Order of Public Reason* does not hold that this question of moral authority is “as problematic and in need of justification as is one’s deployment

of billy clubs” if this means problematic in anything like the same sense. It is a different and more basic issue, viz., whether one’s deployment of billy clubs is to be guided simply by one’s independent moral judgment, or whether we can arrive at commonly recognized authoritative guidelines for when and how billy clubs are to be deployed.

3. WHAT NEEDS TO BE PUBLICLY JUSTIFIED?

The core claim of *The Order of Public Reason* is that these shared rules or guidelines must be justified to all those competent moral agents whose social interactions they structure. If some social-moral rule is not justified to some person in the sense that she has reason to endorse it given her own evaluative standards, then it cannot perform its social normative function with respect to her. When confronted with a demand based on such a guideline, she does not have sufficient reasons to view it as normative; nor does she have good reasons to internalize the rule in the sense of guiding her own action, and producing guilt when she fails to conform. She will, perhaps, accept that it is empirically expected of her to conform, and that others have normative expectations of her, but her own judgment justifiably can deny normative standing to the claim. Recognizing all this — and recognizing that we cannot actually give her adequate reasons to conform — our own reactive attitudes will be undermined. If I say to Betty “You must do this, though of course given your view, you cannot be expected to see why,” and she fails to do it, I cannot see her action as manifesting ill-will to me. I have admitted that she cannot see that what I am on about gives her any reasons, nor do I think it is reasonable to expect her to see such reasons. So where is her ill-will? And if there is no ill-will, the Strawsonian reactive attitudes do not get a grip.

So to say that a publicly justified social morality treats all as free and equal is, essentially, to say just this: that given her own view of normativity, each endorses the rule. Because each has her own reasons to endorse the rule's normativity, no one is merely be told to act as others say they must. Unless the rules of social morality can be so justified, they fail to do the jobs we need them for. Mack, I believe, sees this, but he sometimes reads the book as if were an account of moral truth or moral rightness and wrongness *per se*. Consequently, he tends to merge the public justification requirement concerning social morality with two quite different ones. For example, he reads *The Order of Public Reason* as relying on a "meta-rule that condemns treating others in ways that cannot be justified to them" (Mack ms. p. 27). Now some Kantian-inclined philosophers do indeed advocate such a rule, but I am explicit in distancing myself from it and from them (*OPR*, pp. 17-20). What needs to be justified are the rules of our social morality that purport to supply us with a common, social, normative framework, for without such justification these rules cannot fully perform the functions that are their *raison d'être*. But *The Order of Public Reason* offers no reason why I cannot be treated in ways that cannot be justified to me. When I go to a bar and see others buying beer for strangers but not for me, this beer distribution may not be justifiable to me, but nothing in *The Order of Public Reason* implies that my status as a free and equal person has been insulted. The core supposition is that, if we seek a social morality that fulfills its functions and grounds our reactive attitudes, this morality must not make moral demands on a person that cannot be justified to her.

Neither is it the case that a moral judgment that someone has done wrong must be justified to that person (Mack, ms. p. 27). Again, we must distinguish one's own individual judgments from the shared practice of interpersonal claims, demands, and responses. It is important to keep in mind that the shared normative framework

does not supplant one's individual moral judgment based on one's understanding of the moral truth or whatever one takes to be morally ultimate. Exercising one's individual judgment, one may conclude that, say, allowing the use of billy clubs to kill turkeys for dinner is not in accord with the moral truth; one may be convinced that on the true morality, this is not be allowed. One may condemn such clubbing, and believe it to be wrong. This is all intrinsic to one's moral reasoning, and it cannot be the case that one must publicly justify one's moral conclusions. On the other hand, to advance prescriptions backed up by the rational reactive attitudes, to appeal to rules that one has reasonable grounds for thinking that others have good reason to internalize and live by, to make demands that one reasonably expects that will have uptake by others, and to live with others on terms that do not require them to conform to your individual deliberations and an implied claim to superior moral insight — all this requires public justification.

4. THE OPTIMAL AND THE ELIGIBLE (AND TWO WAYS TO DETERMINE THE ELIGIBLE)

As I remarked at the outset, the dominant view in contemporary social and political philosophy locates true moral normativity in one and only one place: what is morally true, or best, or just. Each is to do her best to decide what that is, and having done so is morally required to act in accordance with her judgment. In the case above, if Alf believes that it is wrong to billy club turkeys, how can he participate in a social moral framework that allows it? This seems paradoxical.

The Order of Public Reason does not dispute that, for each person, she only has one source of normative judgment — her overall set of evaluative standards. In contrast to Kurt Baier, who insists that there are both self-anchored and society-anchored normative guidelines, I do not bifurcate a person's normative considerations in this

way; nor do I bifurcate them into what is shared with others and what is not.¹¹ But I do insist that we distinguish the *optimal* from the acceptable or *eligible*. Suppose one is evaluating a set of proposed moral rules to regulate the use of billy clubs. One's evaluative standards will identify one rule as best (or tied with others as best); if one were the moral dictator or supreme leader, this would be society's rule. But each of us is only one among many and we seek a shared normative framework; we are cognizant from both the objective and subjective perspectives of the great goods that follow from such a shared normative framework. So in *The Order of Public Reason* each Member of the Public asks herself two questions. (1) How would I rank the proposals of the other participants ($p_1 \dots p_n$) in the practice about what our rule (say, about using billy clubs) should be? (2) Given my evaluative standards and my appreciation of the importance of a shared social-moral framework, at what point would I say that living under some proposal p_i is no better than having no shared normative framework at all over this issue? (Of course I would still have my own moral views.) Some proposals may be so misaligned with my evaluative standards that I simply cannot accept a practice structured by them as truly normative. However, the crucial claim of *The Order of Public Reason* is that a person's eligible set is not equivalent to her optimal choice. For example, Mack is a well-known Lockean libertarian; those acquainted with his writings have a good idea of his optimal rules. But in his life, in his participation in market relations, publishing in journals and so on, he employs and relies on social normative standards that depart from what would pertain under his ideal, but do not so drastically depart that he cannot view them as normative, issue complaints grounded on them, feel resentment and indignation when others violate them, and so on.

One of the claims made in the book is that, within the eligible set of each are what I call "jurisdictional rights" — on some matters we handle disagreements in

evaluative standards by allowing each some social sphere in which her standards hold sway. Property and privacy rights are quintessential instance of jurisdictional rights. Mack is unconvinced:

“[A] For, when individuals are allowed to consult their respective full and divergent evaluative standards, it seems likely that reasonable individuals who attach very high value to comradely collective decision-making will strongly prefer jurisdictional rights that are markedly less strong and extensive over the robust commitment to jurisdictional rights that is endorsed at the level of agency. (They will be willing to risk non-rational or inconsistent collective decisions in order to preserve more opportunity for joint decision and joint action.) [B] It seems that this reasonable preference of some individuals is as apt to upset the apple cart of strong and extensive jurisdictional rights at the testing for stability stage as some people’s reasonable opposition to required assistance to the undeserving is apt to upset the apple cart of strong and extensive rights to assistance at the testing for stability stage. “ (Mack ms. p. 19)

Consider first part A of the passage. Mack’s worry is that, rather than endorsing jurisdictional rights to disperse authority (say, over what shall be preached in a church), some may “strongly prefer” strong collective rights according to which we democratically decide on what religions are to be preached. We can think of this as some strongly preferring Constant’s liberties of the ancients.¹² The analysis of *The Order of Public Reason* certainly allows this.¹³ I do not deny that *some* will have a strong preference for the collective decision making of the ancients; what I deny is that *all* will have this option in their eligible sets (e.g., those who are sure that they

will be a minority and will be unable to practice their religion). Figure 1 gives a stylized ordering depicting this situation:

<i>Alf</i>	<i>Betty</i>	<i>Charlie</i>	<i>Doris</i>
Only Catholicism preached	Collective decision-making	Only Lutheranism preached	<i>Jurisdictional Rights</i>
Collective decision-making	Only Catholicism preached	<i>Jurisdictional Rights</i>	Only Lutheranism preached
<i>Jurisdictional Rights</i>	<i>Jurisdictional rights</i>	NO RULE ABOUT THIS ISSUE	Collective decision-making
NO RULE ABOUT THIS ISSUE	NO RULE ABOUT THIS ISSUE	Collective decision-making	NO RULE ABOUT THIS ISSUE
Only Lutheranism preached	Only Lutheranism preached	Only Catholicism preached	Only Catholicism preached

Figure 1

Of course one can still insist that Betty, who “strongly prefers” collective decision-making, would prefer no rule at all to individual jurisdictions, but this begins to look rather more argumentative. After all, individual jurisdictions more or less arose in Western Europe because many thought them second or third best options, while others though them truly best. They really are an effective way for us to live together on terms all can accept; both Lutherans and Catholics have places where their beliefs structure social life. The key to *The Order of Public Reason* is to distinguish what all can live with from why some badly want, but which others cannot see as acceptable.

In this case, then, although some individuals have an option (collective decision-making about what is to be preached in churches) in their eligible sets, that option is not included in the socially eligible set (those options that everyone endorses as sufficiently normative). Notice that this conclusion involves certain restrictions on what a Member of the Public’s eligible set can be; it allows for a great diversity in

orderings, but advances two claims: (1) all will have jurisdictional rights somewhere in their eligible sets and (2) collective decision making will not be in everyone's eligible set. (If (2) is false some societies could justifiably have such collective decision-making; I'd rather doubt this, but it would not be a cause of great concern.) In any event, given these two restrictions, the result is obtained by appeal to the social choice mechanism (essentially a Pareto rule, such that all must prefer¹⁴ a rule to no rule at all), which eliminates the collective decision-making proposal. Let us call this the *Paretean argument via restriction*.

In Part B of the above passage Mack refers to a different argument that advances a stronger restriction on each individual's orderings. Let us say that a *specification of all eligible viewpoints* hold that for some issues, all Members of the Public would essentially agree on the same proposal. In particular, I employ such an argument when maintaining that given each Member of the Public's conception of herself as an agent devoted to the pursuit of her own evaluative standards, each must insist on rights that protect her agency from attempts by others to subvert it. Arguments that involve *specification of all eligible viewpoints* have the strongest assumptions of all the arguments I employ, for they are nothing less than a specific idea of the sort of persons whom we can see as intelligible partners in our social morality, and how one must think of oneself. Do we really agree on agency rights? Well, consider someone who has no deep concern for the protection of his own agency from forceful interventions and from being harmed in his social interactions. He hangs a sign on himself: "To all strangers: Hit me, harm me, force me to obey — use billy clubs if you wish. I advance no claims on these matters." Such people are the stuff of philosopher's fantasies, but those who proclaim this are not recognizable as normal agents living a life with us (and so they are apt to end up in psychiatric wards). What reasonable person disavows a right not to be harmed willy-nilly by strangers?

Now just as agents qua agents must insist on freedom and claims against being harmed, it would seem that qua agents they must insist on assistance when their agency gets in trouble. However, Mack correctly notes that I *reject* this stronger specification of all eligible viewpoints, *viz.*, that all Members of the Public must insist on strong welfare rights. Although each of us may have some reason to welcome such rights (since they benefit one's agency), I do not believe that endorsing such rights is a precondition for being an intelligible agent in a social morality. Such endorsement is not at the very foundation of our understanding of the nature of the others with whom we are cooperating. Someone may quite intelligibly argue that a person should only qualify for assistance if she deserves it, so on her view there is no unqualified right to assistance. Such a person does not seem unintelligible as a fellow participant in social cooperation. She hangs on herself a sign: "To all strangers: No underserved help wanted or given." Such signs are not a fantasy at all. We can understand a social life with people who take that view.

Yet Mack presses: if a person can thus reason to renouncing welfare rights, "why not similarly think" that, once a person tallies up all her values, she will endorse the claim that "people's claims to liberty are to be honored only if those people deserve liberty?" (Mack ms, p. 17). I have given one answer based on the intelligibility of the person who rejects such a claim. But let us allow that there are people who insist that, say, protection from harm should be deserved (they do not seem to be a vocal minority, but let us do philosophy and allow that "someone could say" this). If so pressed we might abandon the more demanding argument from *specification of all eligible viewpoints*, which holds that we all take pretty much the same view of liberty rights, and resort to a *Paretean argument via restriction*. Thus we might have Figure 2.

<i>Alf</i>	<i>Betty</i>	<i>Charlie</i>
Liberty/Harm rights for all	Liberty/Harm rights for all	Liberty/Harm rights for those who deserve it
Liberty/Harm rights for those who deserve it	NO RULE ABOUT THIS ISSUE	Liberty/Harm rights for all
NO RULE ABOUT THIS ISSUE	Liberty/Harm rights for those who "deserve" it	NO RULE ABOUT THIS ISSUE

Figure 2

Figure 2 supposes the same two restrictions as in Figure 1. (1) Liberty rights for all will be in everyone's eligible set, since unless one has liberty rights and claims against harm one cannot really participate in social life (only crazy people wish to hang *that* sign on themselves). Drop liberty and harm rules, and we drop the entire practice. (2) Those like Betty will see that liberty rights for those that "deserve" them is the same as no liberty at all for her, so she will not see why she should endorse as normative a scheme that excludes her and others from its core benefits. I simply do not see how claim (1) can be advanced for strong rights to assistance; we know that large parts of our population *do* reject these rights as unwarranted and we do not wonder whether they are crazy. (As I point out in the book, those who believe that people ought to take responsibility for their own welfare are distributed across socio-economic classes; European evidence indicates there is only a modest correlation of views about individual responsibility for welfare provision with income and class; *OPR*, 364).

That strong rights to assistance cannot be justified as something like first principles of social morality does not imply that they are not politically justified. As I

argue later in the book, once we consider the justification of systems of property rights in the political sphere, claims to a fair share in the fruits of the system of property reemerge as *the fundamental issue in politics*. In *The Order of Public Reason* the justification of welfare rights is a matter for political resolution, part of our larger disputes about what system of property rights we shall have, not fundamental to the very idea of being an intelligible agent. In Rawlsian terms, while agency rights are settled by basic justice and are constitutional essentials, welfare rights are matters for legislative politics. And that does seem to be the way we in fact divide these issues up. So the analysis of *The Order of Public Reason* seems to track the public perception of the different spheres in which these issues arise.

Contemporary political philosophy is often obsessed by the “I can deny that” or “someone can hold...” “counterexample.” And just about anything can be, and has been, denied in philosophical debate, including belief in the external world. Yet when we look around us we see that there is no real debate about the fundamental place of liberty and harm rules in western democracies (though we continue to debate their contours), while there is protracted, deep, and fundamental debate about the justification of assistance owed to others. Some things have been long off the agenda because we all endorse them; they now go the core of how we think of ourselves. To say that they are as disputable as the lightning-rod issues of politics strikes me as just philosophers talking.

5. THE BOUNDS OF THE PUBLIC

It should be clear that the results we can derive from the argumentative machinery of *The Order of Public Reason* are sensitive to the set of evaluative standards that we attribute to the Members of the Public. Mack expresses some concern:

the Deliberative Model includes a constraint of “mutual intelligibility” ([OPR]279) that seems to filter out much more than truly unintelligible proposals. For this constraint disallows the participation in public deliberation of “monomaniacs... such as one who cares about nothing but counting blades of grass, or his stamp collection.” ([OPR]281) The constraint also disallows the participation of those with “evaluative standards that disvalue the very idea of morality, value immoral acts qua immoral acts ...” ([OPR]280) and so on. Now Gaus seems correct to say that there is something unintelligible about persons who have these deviant standards *entering into deliberations for the sake of identifying a justified social morality*. Nevertheless, there seems to be a danger here that success in the public justification of a social morality is being *assured* by banning from the deliberative process any agent whose evaluative standards promise to thwart that success (Mack ms. p. 13).

Our question is this: in a deeply diverse and wide-ranging society, how great is the admissible diversity that will still allow us to maintain a free and stable social moral framework? *The Order of Public Reason* seeks to model how wide-ranging that diversity might be. I suggest an iterative procedure: first see how much we can press the bounds of inclusiveness, and then evaluate the results (*OPR*, 282). That is what I do in the book. If we get some results with something pressing the bounds of inclusiveness, all is well. If I have been too restrictive, and we can relax some of our specifications of the nature of the evaluative standards of the Members of the Public and yet achieve a robust social morality, all the better. If, however, we find that with the restrictions I have supposed we still cannot justify a social morality that all can live with, we must admit that we cannot accommodate a social-moral life with as

diverse an array of valuational perspectives as we hoped. This would be a great cost, but if that is our conclusion, we must see how to cope with it.

The point is that there is no definitive way to set the boundaries of the model. We are searching for a social-moral framework that can secure the benefits of a free social life among as diverse an array of valuational points of view as possible. Of course there are limits, and of course we must suppose some specifications. But we begin by looking around us, and seeing a wide diversity of good-willed and fairly reasonable people with whom we would wish to share a free social-moral life (one not founded on Alf simply making Betty live according to his standards). We ask how might we identify such a free common life? We begin with the Deliberative Model, identifying what you and I take as the features of this population of good-willed and reasonable people, and see what sort of shared social morality they can accept. Our assumptions are always subject to revision, and Mack is entirely right that if we find that inclusion of some evaluative standards make such a life impossible, we may well have to confront the possibility that we cannot share a free social life with them, and consider the costs of that. Given that we are not seeking to track an external truth, but find a way through our pressing problem of a free social-moral life under conditions of deep pluralism, we have no other choice.

6. ON FOLLOWING THE RULES OF SOCIAL MORALITY

I am especially grateful to Mack for raising some deep and important issues about the account of rule-following advanced in *The Order of Public Reason*. The analysis of rule following is fundamental to the enterprise, and thus far few commentators have paid much attention to it.

Contemporary analyses of rule-following divide as to whether rules are best seen as constraints on reasoning or objectives in reasoning.¹⁵ Mack is known for his advocacy of the constraint (or restraint) view.¹⁶ Roughly, on this view a rule prunes one's decision tree; options that previously were open to you are removed; a rule eliminates a decision node by prohibiting or requiring an action. On the alternative view, norms are arguments in a utility function: something that a person sets out to secure.¹⁷ To a considerable extent how we regard rules is a modeling decision: I believe that modeling rules as preferences rather than constraints has been productive in modeling and is consistent with empirical research about the ways people respond to norms. In particular, people's rule-governed actions do indeed seem cost-sensitive, and the rules-as-objectives view readily models this. As with any other objective, on the rules-as-objectives view normative action is sensitive to opportunity costs; if we make it costly to obey, or costly to cheat, we will expect a change in people's behavior.

Mack recognizes this feature of the account, but worries:

Gaus does not take a reasonable agent's concern for rules to be absolute and he provides an extended discussion of how agents (rationally?) weigh values and ends against commitments to rules. ([OPR]148-162) However, I suspect that, if goal-pursuit and rule-following can be as readily balanced as Gaus suggests, then the deontic character of those rules will be undermined or at least those rules will be rendered insufficiently stringent to play their authoritative role (Mack ms. 9-10).

Balancing metaphors can easily lead us astray. That one has some trade-off rate between x and y does not imply that one "balances" x and y in the sense as seeing them as something like equal: it may well be the case that one always chooses x . But

the fundamental point is that our evidence indicates that people's conformity to rules simply is cost sensitive. It will do no good to advocate an account of social-moral rules that claims that they are absolute, as if that settles the matter. Because reasonable people do consider the opportunity costs of compliance, an effective social morality is one that is internalized by the population in such a way that the overwhelming majority is typically willing to incur very significant costs to their personal goals to act on the rules. This is a social and moral achievement, not a matter of conceptual analysis. But there is no reason to think that this willingness to incur costs must be uniform across rules, or even across the population.

This social and moral achievement, I argue, is secured when people come to care about the social-moral rules of their social order. This worries Mack even more:

A related worry is that, because of the centrality of agent's *caring* about rule-compliance within Gaus' own view, a criticism that Gaus offers of the instrumentalist stance may in the end also apply to Gaus' position. The criticism is that, "In the end, the instrumentalist only 'owes himself' fidelity to the rules of morality; even if it would be better for him to owe it to others, he has no reason available to him to recognize their standing ([OPR]186)." If rules come into one's evaluative standards by way of one's *caring* about one's compliance with them, then they may not provide the right sort of reason for rule compliance. For they may only provide reasons having to do with one's own satisfactions and not reasons that derive from others' status as free and equal agents (Mack ms. p. 10).

This is an important and subtle point. The instrumentalist, I argue, has the wrong sorts of reasons to comply with moral rules. A moral rule is justified to the instrumentalist only on the grounds that complying with it satisfies her values. Thus, when she is contemplating violating a rule, the only reason we can give her to

comply is that so doing will advance her own values and projects — that compliance is the prudent option. But a failure to act morally is not a failure of prudence. When Alf complains that Betty has wrongfully used a billy club on him, his complaint is not that she has set back her values, but that she has harmed *him*. Now Mack presses: if Betty puts weight on the harm rule because she cares for it, doesn't the same problem occur? Isn't Alf reduced to reminding her that she cares for the rule, and that is why she should act on it?

There is a sense in which *anything* you can say to me that will have uptake on my part about why I should give you x , which is your due, will be about me — it will be about *my* reasons, *my* values, *my* carings, and so on. However, just because I have an attitude towards x , it does not follow that my reason to x is to satisfy that attitude. Everything depends on where the attitude comes in — whether it is part of the content of my reason, or whether it is an attitude towards a content. Contrast the two intentional states:

α (Pro-attitude [my goal G]) \rightarrow (Pro-attitude [x is due to Eric]) because this will advance G .

β (Pro attitude [Norm N < x due to Eric>])

Now α is indeed the wrong sort of content: my aim is to advance something I care about — G — and so giving x to Eric is a means for me to secure G . However, that is not an objection to β , even though it involves an attitude. The attitude in β is towards the norm, and so not part of the content of my reason, whereas it is part of the content of my reason in α . Of course in both cases the attitude is mine, but I am not aiming at satisfying an attitude in β , whereas I am in α .

The important point is that when you act what you care for, you are not always acting for yourself. I care — a great deal — for my daughter. Indeed, I love her. And any explanation of why I will set aside my own interests for hers, the fact that I love her will be central. We should not give in to a misguided rationalism that would hold that, since I love her so much, I am not really acting for her, but for my own love, and so myself, when I act for her. We must distinguish the content of our attitudes if we are to distinguish selfish from benevolent and pro-social ones.

Although caring is an affective state, we should not suppose that carings are unrelated to reasons — as if we typically care about things for no reason at all (cf. Mack ms. p. 9). When we see that something is hostile to our values and concerns, our caring is apt to be undermined. A main theme of *The Order of Public Reason* (especially chapter IV), is the complex interplay between emotion and reason in our moral life.

7. CONCLUSION

Like Darwin, I “fully subscribe to the judgment of those writers who maintain that of all the differences between man and the lower animals, the moral sense is by far the most important.”¹⁸ As Darwin stressed, though, this is a socialized moral sense, and it is embodied in the social morality that makes human society possible. Although social morality is the foundation of the eusocial life of humanity, it has become almost invisible to moral and political philosophy. Instead, moral and political theory has witnessed complex and controversial constructions expressing the individual moral sense of particular philosophers, with little recognition that it is the coordination of moral sensibility that makes a free social life possible. *The Order of Public Reason* seeks to explore this social morality from the outside, but most

especially from the inside, trying to understand its importance for our lives, and how a non-oppressive social morality requires us to seek a convergence of our moral sensibility — and, hopefully, how this can be achieved. If I have helped convince some that this is a project to be pursued, I will count the book a great success.

Notes

¹ Fred D'Agostino focuses on this feature of *The Order of Public Reason* in his "The Orders of Public Reason," *Analytical Philosophy*, forthcoming.

² Mack writes: "However, unlike Mill, Gaus thinks that we need *authoritative* (and non-political) moral rules and that such authoritative rules can be publically justified. Hence, in a departure from Mill that I do not think he notes, Gaus holds that some of what Mill called 'moral coercion' is publically justified." (Mack ms. p. 2) The implication here is that Mill does not think that some moral coercion can be so justified; but recall that Mill writes: "The object of this Essay is to assert one very simple principle, as entitled to govern absolutely the dealings of society with the individual in the way of compulsion and control, whether the means used be physical force in the form of legal penalties, or the moral coercion of public opinion. That principle is, that the sole end for which mankind are warranted, individually or collectively, in interfering with the liberty of action of any of their number, is self-protection." [John Stuart Mill, *On Liberty*. In *The Collected Works of John Stuart Mill*, edited by J. M. Robson, Toronto: University of Toronto Press, 1963: vol. 18: p. 223. Emphasis added.] Certainly the "moral coercion of public opinion" can be justified regarding actions that concern others.

³ See Kurt Baier, *The Rational and the Moral Order: The Social Roots of Reason and Morality* (La Salle, IL: Open Court, 1995); P.F. Strawson, "Social Morality and Individual Ideal," *Philosophy*, vol. 36 (1961): 1–17.

⁴ I have argued for this interpretation in "On the Appropriate Mode of Justifying a Public Moral Constitution," *The Harvard Review of Philosophy*, forthcoming.

⁵ I am thinking here, obviously, of Joseph Raz's service conception of authority. See his *Morality of Freedom* (Oxford: Clarendon Press, 1986), Part I.

⁶ See Thomas Christiano, *The Constitution of Equality: Democratic Authority and its Limits* (Oxford: Oxford University Press, 2008).

⁷ See my "Public Reason Liberalism" in *The Cambridge Companion to Liberalism*, edited by Chandran Kukathas and Stephen Wall (Cambridge: Cambridge University Press, forthcoming).

⁸ See here *Public Reason in the History of Political Philosophy*, edited by Piers Turner and Gerald Gaus (New York: Routledge, forthcoming).

⁹ Especially important has been the work of Cristina Bicchieri, in particular her *Grammar of Society* (Cambridge: Cambridge University Press, 2006) and and The Bloomington School, led by Elinor

Ostrom. On the latter see Paul Dragos and Peter J. Boettke, *Challenging Institutional Analysis and Development: the Bloomington School* (New York: Routledge, 2009).

¹⁰ Immanuel Kant, *The Metaphysical Elements of Justice*, 2nd edition, edited and translated by John Ladd (Indianapolis: Hackett, 1999), p.116 [§43]. Emphasis added.

¹¹ Baier, *The Rational and the Moral Order*, chap. 5. On rejecting the sharing view see my essay, “A Tale of Two Sets: Public Reason in Equilibrium.” *Public Affairs Quarterly*, vol. 25 (October 2011): 305-325.

¹² Benjamin Constant, “The Liberty of the Ancients Compared with that of the Moderns” in *Political Writings of Benjamin Constant*, edited by Biancamaria Fontana (Cambridge: Cambridge University Press, 1988): 308–28.

¹³ Admittedly, I was more explicit about this in “The Demands of Impartiality and the Evolution of Morality” in *Partiality and Impartiality: Morality, Special Relationships, and the Wider World*, edited by Brian Feltham and John Cottingham (Oxford: Oxford University Press, 2010): 42-64.

¹⁴ To call these “preferences” is not to say that they are desires or mere likings, but they are pairwise orderings in terms of choice worthiness. A preference is simply a pairwise ordering relation—it is not the basis of that ordering.

¹⁵ I have greatly benefited from discussing these issues with John Thrasher and reading his paper “Rule-Following: Constraints and Objectives,” delivered to the *Southern Economics Association*, Washington D.C. (2011) .

¹⁶ See his “Moral Individualism: Agent-Relativity and Deontic Restraints,” *Social Philosophy and Policy*, vol. 7 (September 1989): 81-111. For more recent and formal analysis employing this approach, see Joseph Heath, *Following the Rules: Practical Reasoning and Deontic Constraint* (Oxford: Oxford University Press, 2008); Amartya Sen, “Maximization and the Act of Choice” in his *Rationality and Freedom* (Cambridge, MA: Harvard University Press, 2002): 159–205.

¹⁷ For advocates of this view, see Bicchieri, *The Grammar of Society*, chap. 1; Herbert Gintis, *The Bounds of Reason: Game Theory and the Unification of the Behavioral Sciences* (Princeton: Princeton University Press, 2009), p. 233; Samuel Bowles and Herbert Gintis, *The Cooperative Species: Human Reciprocity and its Evolution* (Princeton: Princeton University Press, 2011), p. 4, chap. 10.

¹⁸ Charles Darwin, *The Descent of Man*, second edition (New York: Penguin, [1879] 2004), p. 120.