Rationality, reasons, and rules

15 Rationality, reasons, and rules

Brad Hooker

Introduction

I had already started learning from Hanjo Glock before I joined him at Reading in 1993. Once I became his colleague, the rate at which I absorbed from him philosophical arguments and insights and wider wisdom was humbling. And the example he set—of ambition, energy, conscientiousness, fortitude, wit, and good humour—inspired everyone around him. I was in awe of him then, and this awe has increased in the years since, as his arguments have gained adherents across many areas of philosophy.

Rationality

Glock writes,

> Within contemporary academic debates one can distinguish four general conceptions of rationality. According to the first, it is the capacity to maximise satisfaction of one's interests or goals; according to the second, it is responsiveness to reasons; according to the third, it is the ability to reason—draw theoretical and/or practical inferences and to avoid inconsistencies; according to the fourth, it is the ability to justify one's actions and beliefs to others.
>
> (2019a, 665)

I agree that these four general conceptions of rationality are especially prominent. And I think Glock and I agree which of these four are most plausible.

Before I start on explaining why, let me make a point of clarification about what I think Glock meant by 'conceptions of rationality'. I think he meant competing conceptions of what constitutes rationality. *Having rationality* is compatible with *being irrational*. Having rationality is a matter of possessing capacities, and being rational is a matter of correctly exercising those capacities.

DOI: 10.4324/9781003202929-20

The first conception of rationality that Glock lists is the idea that rationality is the capacity to maximise satisfaction of one's interests or goals. This idea is not as simple as it initially looks. Admittedly, there are simple cases. For example, suppose your chief desire right now is for a cold drink and so you go now to the refrigerator to take water from the pitcher there. In other cases, however, you have a desire that needs *further specification* before you start thinking about means to satisfy it. For example, you want to achieve something important this week, but you need to think about whether finding a merely temporary solution to a problem afflicting your community qualifies as an important achievement (cf. Williams 1981, 104). Another example might be that you want to find something that appropriately symbolises your friendship with Kayta, but you have to think about whether a beautiful bowl or a long-life house plant would be better.

The idea that rationality consists in the *capacity* to maximise satisfaction of one's interests or goals is often associated with the idea that rationality *requires* one to exercise this capacity if one has the capacity. Suppose my *sole* goal right now is to get home by the quickest route and, fully aware of what I am doing, I turn right, despite knowing that turning right will not get me home by the quickest route. The previous sentence might strike you as somehow inaccurate. You might surmise, for example, that getting home by the quickest route must not have been my only goal, or that I must have failed to appreciate that turning right would take me away from the quickest route. However, if you accepted that I really did have only one goal and that I really did appreciate that turning right would not achieve that goal, then you would think that I was irrational, indeed that what I did was barely intelligible, given my mental states at the time. How could I really have getting home by the shortest route as my goal, know that turning right would not achieve that goal, and yet turn right in order to achieve the goal?

Actually, people sometimes do have a goal, know that a certain kind of action will not achieve the goal, and yet choose an action of that kind in order to achieve the goal. For example, I want to impress you, I know that bragging will not impress you, and yet I find myself bragging in order to impress you. Sometimes I behave irrationally. Sometimes other people do too.

In many situations, one does not know in advance what all the consequences of this or that act would turn out to be. Thus, in many cases, one does not know in advance what the most efficient means to one's goal would be. In such cases, rationality cannot reasonably insist that one choose the most efficient means, since one does not know which means would be most efficient.

In cases of uncertainty, one often has *some* information about the *probabilities* of outcomes of possible actions. If the information I have suggests that turning left will very probably get me home sooner than

not turning left, then, as long as my sole aim is to get home as quickly as possible, rationality requires me to turn left. More generally, when facing a choice between actions whose *actual* consequences one cannot know in advance, one would be irrational to ignore evidence about *likely* consequences (unless one has grounds to think this evidence is misleading). A way of incorporating probabilities into one's account of rationality is to conceive of rationality as combining judgements about the values of achieving one's different goals and the probabilities that alternative actions would achieve those goals.

Now we should turn to the question of whether rationality requires one to prioritise one's own interests and goals over other people's. The idea that rationality consists in the capacity to maximise satisfaction of one's interests or goals does *not* dictate that one's goals must be focused on one's own life. Whether your goal is to support Amnesty International or the United Nations International Children's Emergency Fund or your impoverished cousin, rationality conceived of as the capacity to maximise satisfaction of one's interests or goals does not denigrate your altruistic goals. Indeed, this conception of rationality is completely non-judgemental about your goals and indeed about everyone else's.

While this conception of rationality does not denigrate altruistic goals, it also does not insist one have altruistic goals or concerns. My own view is that there is something rationally defective about agents who are concerned exclusively with their own goals and attach no non-instrumental importance to other people's worthwhile goals.

But a reply to my view might be that the concept of individual agency draws a sharp line between the agent's goals and intentions and other people's goals and intentions, and that thus *rational* individual agency cannot discard this sharp line. I agree that the concept of individual agency would be severely threatened if each agent were rationally required to give the goals of every other agent the same weight as his or her own. But my much more modest contention is that rational agents must attach *some* non-instrumental importance to other people's worthwhile goals. This modest contention hardly threatens the concept of agency.

I admit that taking the distinction between one's own goals and those of others to be practically important is not arbitrary, even if this distinction can be taken too far. However, there are other possible distinctions that are, beyond question, arbitrary. One is Derek Parfit's example of someone who cares about benefits or harms to himself that occur on any day of the week except Tuesday (Parfit 1984, 124–126). Here, someone singles out an arbitrary period of time—Tuesdays. Someone else might single out an arbitrary space. She might, for example, care about what benefits or harms herself anywhere except in Winnipeg.

Both the time and space kind of examples have an infinite number of instances. Just as it is irrational to care about what happens to oneself *except on* Tuesday, it is irrational to care about what happens to oneself

except on Saturday, or *except at* two minutes to the hour, etc. Likewise, it is irrational to care about what happens to oneself *only on* the 15th of the month. And just as it is irrational to care about what happens to oneself *except in* Winnipeg, it is irrational to care about what happens to oneself *except in* some other particular place, or *only in* some other particular place.

In response to these kinds of examples, we might point out that the people in these examples do not deny that they have interests and goals on Tuesdays, just as they have interests and goals on other days of the week, and at two minutes to the hour, and at any other particular time, and when they are in Winnipeg. Hence, if rationality requires one to maximise the satisfaction of one's interests and goals, then the satisfaction of one's on-Tuesday interests and goals and one's in-Winnipeg interests and goals *are* to be counted in the aggregate of one's interests and goals to be maximised. And if rationality can require one to care about, or at least count in one's calculation of what is to be maximised, one's future interests and goals even though one does not *now* care about them, it is not such a big step to holding that rationality can require one to give some weight to the interests and goals of others even though one does not *now* care about them.

I turn now to the second conception of rationality that Glock lists, the conception of rationality as responsiveness to reasons. If we think that people *have reason* to maximise the satisfaction of their interests and goals, then we may well think that *rationality* calls for people to *maximise* satisfaction of their interests and goals. One of the most often discussed examples of irrationality is the pursuit of present benefits when these are known to lead to greater harms later (Sidgwick 1907, 124, n. 1, 381; Nagel 1970, chs. 6–8; Parfit 1984, 133, 158–163). Imagine someone who gives himself a weekend of leisure though he knows he needs to be studying assiduously throughout the weekend in order to pass the test he faces next week. It is very tempting to say that he is irrational in choosing the relatively minor present benefit of a weekend of leisure over the far more significant future benefits that would come from passing the test.

The benefits to him that would come from passing the test provide reasons for him to do what is necessary and sufficient for him to pass the test. He is much more likely to pass the test if he tries to pass it, and if he studies assiduously throughout this weekend for the test. Admittedly, the pleasure he would get from a weekend of leisure would constitute a benefit to him, and the prospect of this benefit generates a reason for devoting the weekend to leisure. However, the weights of reasons deriving from benefits or harms to him presumably correlate with the sizes of the benefits or harms to him, and the benefits to him of passing the test would be much greater than the benefit to him of a pleasurable weekend. Hence, he presumably has stronger reason to spend the weekend preparing for the test than he does to spend the weekend on leisure.

If we conceptualise rationality as responsiveness to reasons, we might say that what is irrational about indifference to benefits or harms to oneself that occur on Tuesday or in Winnipeg is *not* that such indifference *arbitrarily* distinguishes between benefits or harms to oneself on Tuesday and benefits or harms to oneself on other days, or between benefits or harms to oneself when one is in Winnipeg and benefits or harms to oneself when one is in other places. What is irrational about indifference to benefits or harms to oneself that occur on Tuesday or in Winnipeg is instead that such indifference is unresponsive to one's reasons to care about benefits or harms to oneself that occur on Tuesday or in Winnipeg.

Putting off a more elaborate discussion of reasons until the following section, I turn now to the third conception of rationality in Glock's list, rationality as drawing theoretical and/or practical inferences and avoiding inconsistencies. I think Glock and I agree that this conception of rationality is the most widely shared. Having obviously inconsistent beliefs is irrational. Perhaps failures to draw obvious inferences is also irrational. But we go too far if we insist that all failures to draw unobvious inferences are irrational.

Drawing valid inferences can come at a cost—in terms of effort, time, and cognitive or emotional overload. Because of such costs, rationality must concede that drawing further valid inferences can sometimes not be what agents have most practical reason to do. Imagine some agent sitting in an ivory tower teasing out valid inferences while the city around her burns.

I turn now to the fourth conception of rationality in Glock's list, rationality as the ability to justify one's actions and beliefs to others. In the case of each of the first three conceptions of rationality Glock listed—the capacity to maximise satisfaction of one's interests or goals, responsiveness to reasons, and the ability to draw theoretical and/or practical inferences and to avoid inconsistencies—I have mentioned these as capacities but mostly discussed them as requirements, in the form of 'rationality is not only the capacity to [. . .] but also the requirement to [. . .]. '. In the case of the fourth conception in the list, there seems to me no pressure to think of rationality as a requirement as well as a capacity. Admittedly, in many contexts, there are requirements to justify one's actions and beliefs to others. However, these requirements do not come from rationality.

Justifying one's actions and beliefs to others must be understood as laying out considerations and reasoning which others *might* accept as justifying one's actions and beliefs. Justifying one's actions and beliefs to others *need not entail* that all others would *in fact* accept one's justifications. Some people might reject these justifications because they cannot see why one's conclusions follow from one's premises. And some people might deny one or more of one's premises.

Laying out the considerations and reasoning that justify one's actions and beliefs requires thought—and, often, fairly sophisticated thought.

We need not here get caught up in the question of how sophisticated such thought must be. If rationality consists in the ability to justify one's actions and beliefs to others, one implication is that, whatever the level of sophistication that thought must have in order to lay out for others the considerations and reasoning justifying one's actions and beliefs, any being who is incapable of that level of sophisticated thought is incapable of rationality. And it does seem intuitively correct that any being incapable of a threshold level of sophisticated thought lacks rationality.

What does *not* seem intuitively correct is that you cannot be rational unless you have both the capacity and the opportunity to justify *to others* your beliefs and actions. On the face of things, your ability to *lay out* for others your considerations and reasoning requires that you have the ability to *communicate* these considerations and lines of reasoning. Yet your having the ability *to communicate* is not a necessary condition of your being rational, much less of your having rationality. Imagine that you survived a terrible injury that robbed you of the capacity to communicate your mental states. Obviously, your *loss of the ability to communicate* your mental states does not entail that you *lack mental states* or that you *lack abilities to think, draw valid inferences, entertain hypotheses, evaluate arguments, assess actual and possible beliefs and actions, form desires and intentions, etc*. If you have retained these abilities, then you have rational capacities. You are able to think, desire, and intend rationally, even if you cannot communicate to others what you think, intend, etc.

Perhaps the most charitable way to construe the idea that rationality is the capacity to justify one's beliefs and actions to others circumvents problems about the inability to communicate with others, whether that inability comes from one's injuries or from the absence or deafness of others. Rationality as the capacity to justify one's beliefs and actions to others should be construed as the capacity to lay out, at least in one's own mind, the considerations and reasoning that one takes to support one's beliefs and actions. Possible others might accept them if presented with them.

Reasons

I indicated that I would return to the conception of rationality as responsiveness to reasons (an old haunt; see Hooker 1987). This conception of rationality seems to me one of the two most plausible conceptions of rationality, the other being the conception of rationality as drawing inferences and avoiding inconsistencies. Actually, rationality as drawing inferences and avoiding inconsistencies could be thought of as a subset of rationality as responsiveness to reasons (cf. Kiesewetter 2017; Lord 2018). To draw inferences and to avoid inconsistencies are ways of responding to reasons. For example, you have conclusive reason not to believe more than one of any two inconsistent propositions.

Inferences from premises about probabilities and practical inferences are more complicated. Suppose I knew the homemade curry put in front of me was 50% likely to upset my digestion. Suppose I knew both that I was not hungry and that there was another food choice available to me that was equally delicious and nutritious but very unlikely to upset my digestion. Still, maybe I had most reason to eat the curry because I would have offended my boss if I had not (he was the one who made the curry).

This example illustrates an important aspect of reasons for desiring, intending, and acting. That the curry had a 50% likelihood of making me feel unwell was a reason not to eat it. There might have been other reasons not to eat it, such as that, while eating the curry, I would have been likely to spill it on myself. And yet all the reasons against eating the curry might be outweighed by reasons to eat it. (With respect to reasons for belief, of course there can be evidence *for* some conclusion and evidence *against* that same conclusion.)

One of the great breakthroughs in practical philosophy was W. D. Ross's distinction between prima facie duties and all-things-considered duties (Ross 1930, ch. 2). What Ross meant by 'prima facie duty' was not 'duty on first look', as if the duty would turn out not to be a duty on subsequent investigation. What he meant was 'duty in one respect and to an extent'. Many philosophers have thus changed Ross's terminology from 'prima facie' to 'pro tanto', in order to better express what Ross had in mind.

Ross proposed that, when we face a choice between mutually exclusive acts, we are to weigh up the different pro tanto duties we have in the situation and decide which act has the strongest aggregate of pro tanto duty on its side. The act with the strongest aggregate of pro tanto duty on its side is then the act that is our all-things-considered duty in this situation. The idea that there is a plurality of moral pressures, which do not come in a strict hierarchy of importance, and that moral agents thus need good judgement to adjudicate conflicts between these moral pressures was the central idea of Ross's deontological pluralism in the 1930s, though he was hardly the only proponent. And moral pluralism is best expressed using the pro tanto/all-things-considered distinction.

For example, you have a pro tanto duty to help those in need, especially when you can do so at no cost to yourself. Suppose I am in need and you could help me at no cost to yourself. But suppose there is someone else in even *greater* need whom you could instead help at no cost to yourself. Suppose you cannot help us both. You have a pro tanto duty in one respect (need) and to an extent (based on the degree of need) to help me, but you also have a pro tanto duty in the same respect but to a greater extent to help the other person. Thus, your all-things-considered duty in this situation is to help the other person.

The term 'duty' is more often used and heard in the pro tanto sense than in the all-things-considered sense. The example in the previous

paragraph concerned a conflict of *pro tanto* duties. But we could have expressed the example in terms of a conflict of *duties*. Saying that you have a duty to do something in a situation is not to proclaim an all-things-considered moral verdict about what to do in that situation, since you might also have an opposed duty in this situation to do something else. I am not denying, however, that we can also use the word 'duty' to mean an 'all-things-considered moral verdict about what to do'. Indeed, precisely because the word 'duty' is ambiguous between 'pro tanto duty' and 'all-things-considered duty', Ross did the world a favour by distinguishing between these meanings and offering terminology to mark the distinction.

Just as 'duty' has a pro tanto sense and an all-things-considered sense, so does the term 'ought'. However, whereas 'duty' is normally meant and understood to be referring to pro tanto considerations, 'ought' is normally heard as affirming an 'all-thing-considered' judgement. Return to the example where I am in need and you could help me at no cost to yourself, or you could instead help someone else who is in even greater need than I am and your helping that person would impose no cost on you. If someone asserted that you ought to help me, you might reject that assertion on the grounds that there is another person in greater need whom you could also help at no cost to yourself. If others' needs are the determining factor, then it seems natural to say that really what you ought to do is help the other person. I submit that you would *not* be inclined to say that you are subject to conflicting oughts. Rather, you would be inclined to say instead that whatever in the end you should do is what you ought to do, with the outweighed consideration's failing to be an ought at all.

The dominant moral theories in the 1950s and 1960s in Anglophone countries were utilitarianism and Kantianism. Kant's Categorical Imperative tells one what one ought morally to do—act on maxims that one can will to be universal laws; or treat others always as ends in themselves, not merely as means. The Categorical Imperative is not telling one merely what considerations to weigh up when making moral decisions. The simplest form of utilitarianism is also framed as an imperative—do whatever maximises utility, impartially calculated. Both Kant's Categorical Imperative and this simple form of utilitarianism were often framed in terms of 'ought' judgements. These were all-things-considered 'ought' judgements, not merely pro tanto ones.

Monistic moral theories offer a single imperative based on a single consideration, even if that single consideration is complex. The first formulation of Kant's Categorical Imperative, for example, identifies just one thing as determinative, namely whether the agent (arguably Kant meant *every* agent) could will her maxim as a universal law. And the simplest form of utilitarianism is an act-utilitarian theory holding that what one ought to do depends entirely on what maximises aggregate utility, where benefits and harms to everyone, including future people,

are included. Kantianism and this simple utilitarianism are alike in pointing to a single consideration, albeit not the same one. Each of these two theories does not need the pro tanto/all-things-considered distinction because each of these theories denies that there are different kinds of moral considerations to be weighed against one another.

I do not mean to deny that utilitarianism and Kantianism cannot make use of the pro tanto/all-things-considered distinction. For example, even the simplest utilitarianism might allow that each possible benefit generates a pro tanto reason to promote it, and these reasons are to be weighed together and against other such reasons in order to ascertain an all-things-considered moral verdict on what to do. Kantianism might have use of the distinction in other ways. My point was not that the distinction cannot be used by utilitarianism and Kantianism. My point was that act-utilitarianism and Kantianism do not need to use the distinction.

Having commented on Kantian and utilitarian theories, I turn to virtue ethics, by which I mean ethical theories that evaluate action as right or wrong by reference to what a virtuous person would characteristically choose. Does virtue ethics need the pro tanto/all-things-considered distinction?

The answer seems to me to depend on whether the virtues are conceived of as possibly conflicting with one another. One conception is that the virtues cannot conflict, because, for example, a person could not really grasp what kindness requires unless she also understood that kindness cannot require dishonesty. If the virtues cannot conflict with one another, then virtue ethics has no need of the concept of a pro tanto moral reason. In contrast, if the virtues can conflict with one another, then someone might find herself in a situation where, for example, kindness pulls in one direction but honesty pulls in the opposite direction. The pro tanto/all-things-considered distinction helps articulate such conflicts: There is a pro tanto reason to be kind and a pro tanto moral reason to be honest, and what the agent should do, all-things-considered, depends on what a virtuous person in this situation would characteristically take to be the more important moral reason when they conflict in the situation at hand.

During the 1940s, 1950s, and 1960s, most ethical theorists in Anglophone countries directed their attention to utilitarianism and Kantianism, on the assumption that both virtue ethics and moral pluralism are unsatisfactory moral theories. Virtue ethics was presumed to be either implausibly committed to holding that virtues cannot conflict or to be in effect a form of moral pluralism. The prevailing objection to moral pluralism was that it, in D. D. Raphael's words,

> does not meet the needs of a philosophical theory, which should try to show connections and tie things up in a coherent system. To

> look for unity where none exists would, of course, be foolish; if the diversity of moral rules were intractable, it would be pointless to go on searching for some way of tying them up together. But the moral rules of ordinary life are not obviously all different from each other.
> (Raphael 1994, 55)

A related objection to moral pluralism is that it is often less determinate in its practical applications than utilitarianism and Kantianism.

Nevertheless, the prevailing methodology in normative ethical theory since about 1970 seems to me to have elevated moral pluralism to the title of 'theory to beat'. This methodology holds that we should take our considered moral convictions at all levels of generality and try to get them into 'reflective equilibrium' with one another, and with everything else we believe. Thus, we expect our general moral principles to be compatible with our non-moral beliefs and to cohere with the specific moral convictions we have after due consideration.

The methodology of seeking reflective equilibrium is normally traced to John Rawls (1951; 1971, 19–21, 46–51; 1974–1975, sect. 2; 1980, 534). Rawls assumed that the pressure to find general principles is so strong that we should endorse whatever is the most attractive set of general principles we can find and we should discard moral convictions that do not accord with that set. I agree that, if two sets of general moral principles are equally good at supporting our more specific moral convictions but one of the sets of general moral principles provides this support on the basis of fewer principles, this more parsimonious set of principles is better. What makes it better is that it explains equally much on the basis of less. But we should not take the method of reflective equilibrium to be pre-committed to endorsing whatever *single* fundamental moral principle supports the highest percentage of our other moral convictions. For all we know prior to thorough investigation, there is some set of plural fundamental principles that (a) is consistent with our non-moral beliefs, (b) seems plausible, and (c) coheres with all our more specific moral convictions better than any single fundamental principle on its own does.

Indeed, despite the pressure exerted on us by having parsimony and connectedness as desiderata in our moral theorising, we might well think that the method of reflective equilibrium takes us to moral pluralism. Moral pluralism does an unbeatable job of agreeing with our various moral convictions. Unless some single fundamental moral principle can do as good a job of agreeing with our various moral convictions as moral pluralism does, then moral pluralism is the moral theory best justified to us (Hooker 1996; 2000a, ch. 1).

I return to the point that moral pluralism is best expressed using the pro tanto/all-things-considered distinction. The term 'normative reasons' is most naturally understood to be referring to normative favourers or disfavourers, i.e., pro tanto normative considerations rather than all-

things-considered verdicts about what to do. The emphasis on reasons makes sense if we need to acknowledge different, possibly conflicting normative pressures, of possibly different strengths, before we decide what to do. For many of us, everyday practical decision making is sprinkled with episodes of weighing up the strengths of conflicting pressures, leading up to a conclusion about what, all-things-considered, we should do. Thus, thinking in terms of reasons comes very naturally.

~~Section 3:~~ Rules

Trying to think only in terms of reasons and bypass rules entirely, however, would be disempowering. A great deal of thinking involves recognising a difference between things, which often forms the basis of reacting to different things differently. Glock writes, 'Judgement precisely involves classifying an object *as* being of a certain kind. And this idea can in turn be spelled out by saying that judgement requires the deliberate choice between different options in a sorting or discrimination task' (2010, 19). Without wading into the discussion of Glock's arguments about animal thought (which are taken up by other contributors to this volume), I note that, at least in the case of beings with language, to classify or categorise something as being an instance of a kind is a rule-governed activity.

According to Glock, 'The linguistic meaning of expressions depends on general rules. These rules provide standards for the correct use of expressions' (2015, 842). The rules providing standards for the correct use of expressions are *normative* rules shared by the linguistic community. With respect to the rules determining linguistic meaning, I have no reason to dissent from Glock's drawing on H. L. A. Hart (1961):

> [I]n a group G a behavioural regularity R is a *shared rule* if and only if
>
> 1 it is rare for members of G to deviate from R
> 2 if members of G deviate from R, they are subject to sanctions, including the verbal sanction of being criticised
> 3 these sanctions are generally accepted by members of G. (Glock 2019b, 313)

Much of what Glock takes from Hart can be applied to moral rules as well as to rules about linguistic meaning. When we think of possible moral rules, we think of rules to be shared in the sense that

1 People would comply with the rules and at least sometimes use the rules to guide their behaviour.
2 Non-compliance would be met with sanctions, such as blame, indignation, resentment, the withdrawal of good will and cooperation, etc.

3 Both the rules and their connection to sanctions would be accepted as justified.
4 People would think of the rules as generating reasons for action and grounding interpersonal justification.

One difference between linguistic rules and moral rules is that linguistic correctness is determined by linguistic rules that are *already* shared by those with linguistic competence. Already shared rules are hardly static. People can make proposals of linguistic reforms, coin new terms, and innovate in other ways. Nevertheless, linguistic correctness in general is determined by already shared rules.

When we turn from linguistic rules to moral rules, the fact that a rule is already shared seems to me less authoritative. Let us distinguish between a set of rules that are *already* shared and a set of rules that might not be already shared but *ideally* would be shared. There could be a society where the two sets of rules are the same—people there already accept the ideal rules. But the importance of the distinction between established rules and ideal rules comes out when we focus on a society in which the established rules are not ideal. For example, consider a society in which the *actually shared* rules insist that one should unwaveringly identify either as male or as female and then routinely behave very differently depending on which one is. But *ideal* rules would neither demand that one identify as being of either one of two genders nor require different behaviour depending on gender. Because of this point about ideal rules, refusing to identify as being of either one of two genders is morally permissible. Admittedly, there might be powerful self-interested reasons not to offend against the rules accepted by the people around you. But if the rules they accept are unnecessarily restrictive, or invidiously discriminating, or destructively lax, the bare fact that those rules are already accepted does not entail that they dictate moral correctness.

Rules articulating requirements generate reasons for action. For example, there is a moral rule requiring people not to steal, and this moral rule is one source of reasons not to steal. Requirements on action also lead to reasons to react to behaviour with (e.g.) feelings of indignation, resentment, or guilt and withdrawal of goodwill, ostracism, or more (see Scanlon 2013, 105–109; 2018, 120–121).

For rules, including rules making requirements, to be justified, there must be undefeated reasons for having these rules. Reasons for having these rules are more basic than the rules themselves: The reason for having the rules is not that there were already reasons for action that were independent of and prior to the rules. Rather, the reason for having the rules is that people's accepting the rules would have better consequences than people's not accepting it would—where 'better consequences' is not merely a matter of people's more often acting on reasons they already had. To justify these assertions about which reasons are basic, I would

have to digress for longer than tolerable. (But see Copp 2010; 2020; Parfit, 2017, 432.)

Moral rules not only impose requirements but also grant permissions and powers. For example, there is a moral rule permitting people to devote their own time, energy, attention, and other resources to achieving their own goals even when these resources could instead be used to help other people achieve their similar or even somewhat more important goals. Admittedly, this permission is not unlimited. For example, one might need to sacrifice one's goal of getting home early in order to rescue an accident victim whom one comes across on the way home.

The moral rule *permitting* people to devote their own time, energy, attention, and other resources to achieving their own goals does *not* itself give people *reasons* to do what the permission allows. To be sure, people *do* have reasons to devote their own time, energy, attention, and other resources to achieving their own goals. But these reasons do not come from the moral permission they have to do so. The point generalises: Permissions to do things do not *on their own* generate reasons to do what the permissions allow.

Permissions are liberties. If holding on to a particular liberty is important and if occasionally exercising the liberty is necessary to hold on to it, this might give you a reason to exercise the liberty, though not necessarily an undefeated reason. What gives you the reason in such a case is not the liberty *on its own* but rather the combination of the value of preserving the liberty and the necessity to exercise the liberty occasionally in order to preserve the liberty. Nevertheless, a person's permissions and liberties *do* entail *reasons* for action *for other people*. If a person is at liberty to do a kind of act, other people have reason not to force that person not to do that kind of act.

There are also moral rules specifying normative powers (Owens 2012; Hohfeld 2019; Chang 2020). Suppose that up to now you have no right against me that I meet you tomorrow for lunch. But if I now promise you I will meet you tomorrow for lunch, I have thereby created a moral obligation on me and a moral right in you against me that I meet you tomorrow for lunch. Your new right comes with your new moral power to waive your right and cancel my obligation.

The example of promising as a social practice constituted by rules brings out one of the ways in which rules are ineliminable from morality as we know it. (On the kind of rules that constitute a social practice, see Rawls 1955; Glock 2019b, 307.) Many moral actions make sense only when construed as complying with a social practice constituted by rules. To explain what a promise is, we have to refer to rules about agents' voluntarily creating obligations for themselves and rights for others.

Rules constitutive of social practices are not the only ones to pervade moral thinking. Suppose you are trying to decide what to do in a situation in which many innocent people would benefit or be harmed by

your decision. Let us add that some of the alternative acts you could choose would involve varying degrees and kinds of dishonesty. We already know enough about the situation to see that various pro tanto moral reasons come into play. One is a pro tanto moral reason to benefit innocent others, with the more benefit to others and the more others, the better. Another is the pro tanto moral reason not to harm innocent others. And another is the pro tanto moral reason to avoid being dishonest.

While each of these moral reasons is one that *you* have *in this particular situation*, they are instantiations of general phenomena: *All* agents have pro tanto moral reasons to benefit innocent others, not to harm innocent others, and not to be dishonest *in all situations in which acts of these kinds are available*. And such *general* moral reasons are often expressed as rules in the imperative voice—such as the rules 'do good for others', 'don't harm others', and 'do what is honest'. (On the 'inherent generality' of rules, see Glock 2015, 843.)

A point regularly made is that common knowledge that certain rules are widely accepted plays an immensely important role in coordinating people's behaviour. A standard example is that each society needs there to be common knowledge about which side of the road to drive on there. Not every rule solves a coordination problem (Glock 2019b, 312), but many do, and many coordination problems would be very difficult to solve without shared rules.

Perhaps even more importantly, common knowledge that certain moral rules are widely accepted also plays an immensely important role in assuring people about other people's behaviour. Admittedly, we know the law forbids (e.g.) stealing, and law enforcement attempts to deter stealing. Nevertheless, we will be even more confident that others will not steal from us if we believe that, in addition to others' awareness of the law and the punishment imposed for stealing, others accept a moral rule against stealing. What I have written here about stealing also applies to physical aggression and promise breaking. Moral and legal rules had better 'protect persons, property, and promises' (Hart 1961, 193).

Could human societies realistically make do with thinking in terms of reasons and eschew all references to moral rules? I suspect that reducing all moral rules to propositions about reasons could not be done without losing valuable information. However, even if such reduction is possible, I expect that many of the resulting propositions about reasons would need to be long and complicated, thus sacrificing the simplicity that many rules have. Sacrificing that simplicity would have high costs in terms of cognitive overload and the corrosion of assurance (see Hooker 200[illegible]2007).[1]

AU: Please confirm whether the citation of Hooker (2000) should be Hooker (2000a) or Hooker (2000b) or both.

Note

1 In September 2021, this paper was presented at a conference in Hanjo Glock's honour. I thank Christoph Pfisterer, Nicole Rathgeb, and Eva Schmidt for very

helpful comments on a draft given to me prior to the conference. I also thank Maria Alvarez, Gerhard Ernst, Hanjo Glock, John Hyman, Jörg Löschke, and Christoph Pfisterer for comments given to me at the conference. Each of all these people made comments that significantly improved the paper.

References

Chang, R. (2020). Do We Have Normative Powers? *Aristotelian Society* 94: 275–300.

Copp, D. (2010). The Wrong Answer to an Inappropriate Question? In Black, S. and Tiffany, E. (eds.), *Reasons to Be Moral Revisited*. *Canadian Journal of Philosophy* 33: 97–103.

Copp, D. (2020). The Rule Worship and Idealization Objections. *Oxford Studies in Normative Ethics* 10: 131–155.

Glock, H.-J. (2010). Can Animals Judge? *Dialectica* 64: 11–33.

Glock, H.-J. (2015). Meaning and Rule Following. In *International Encyclopedia of the Social & Behavioral Sciences*. 2nd edition. Vol. 14. Oxford: Elsevier, pp. 841–849.

Glock, H.-J. (2019a). Agency, Intelligence and Reason in Animals. *Philosophy* 94: 645–691.

Glock, H.-J. (2019b). Normativity of Meaning Revisited. In Roughley, N. and Bayertz, K. (eds.), *The Normative Animal?* Oxford: Oxford University Press, pp. 295–318.

Hart, H.L.A. (1961). *The Concept of Law*. Oxford: Clarendon Press.

Hohfeld, W. (2019). *Fundamental Legal Conceptions*. Edited by W.W. Cook. New Haven, CT: Yale University Press.

Hooker, B. (1987). Williams's Argument against External Reasons. *Analysis* 47: 42–44.

Hooker, B. (1996). Ross-Style Pluralism versus Rule-Consequentialism. *Mind* 105: 531–552.

Hooker, B. (2000a). *Ideal Code, Real World: A Rule-Consequentialist Theory of Morality*. Oxford: Clarendon Press.

Hooker, B. (2000b). Moral Particularism—Wrong and Bad. In Hooker, B. and Little, M. (eds.), *Moral Particularism*. Oxford: Clarendon Press, pp. 1–22.

Hooker, B. (2007). Moral Particularism and the Real World. In Lance, M., Potrc, M. and Strahovnik, V. (eds.), *Challenging Moral Particularism*. London: Routledge, pp. 12–30.

Kiesewetter, B. (2017). *The Normativity of Rationality*. Oxford: Oxford University Press.

Lord, E. (2018). *The Importance of Being Rational*. Oxford: Oxford University Press.

Nagel, T. (1970). *The Possibility of Altruism*. Oxford: Clarendon Press.

Owens, D. (2012). *Shaping the Normative Landscape*. Oxford: Oxford University Press.

Parfit, D. (1984). *Reasons and Persons*. Oxford: Clarendon Press.

Parfit, D. (2017). *On What Matters*. Vol. 3. Oxford: Oxford University Press.

Raphael, D. (1994). *Moral Philosophy*. Oxford: Oxford University Press.

Rawls, J. (1951). Outline for a Decision Procedure in Ethics. *Philosophical Review* 60: 177–197.

Rawls, J. (1955). Two Concepts of Rules. *Philosophical Review* 64: 3–32.
Rawls, J. (1971). *A Theory of Justice*. Cambridge, MA: Harvard University Press.
Rawls, J. (1974–1975). The Independence of Moral Theory. *Proceedings and Addresses of the American Philosophical Association* 48: 5–22.
Rawls, J. (1980). Kantian Constructivism in Moral Theory. *Journal of Philosophy* 77: 515–572.
Ross, W.D. (1930). *The Right and the Good*. Oxford: Oxford University Press.
Scanlon, T. (2013). Giving Desert Its Due. *Philosophical Explorations* 16: 101–119.
Scanlon, T. (2018). *Why Does Inequality Matter?* Oxford: Oxford University Press.
Sidgwick, H. (1907). *Methods of Ethics*. London: Macmillan.
Williams, B. (1981). Internal and External Reasons. In Williams, B. (ed.), *Moral Luck*. Cambridge: Cambridge University Press, pp. 101–113.