

1

When Is Impartiality Morally Appropriate?

BRAD HOOKER

With respect to morality, the term ‘impartiality’ is used to refer to quite different things. My chapter will focus on three:

1. Impartial application of good (first-order) moral rules
2. Impartial benevolence as *the* direct guide to decisions about what to do
3. Impartial assessment of (first-order) moral rules

What are the relations among these three? Suppose there was just one good (first-order) moral rule, namely, that one should choose whatever one thinks will maximize aggregate good. If there were just this one moral rule, then impartial application of that one rule might be compatible with impartial benevolence as *the* direct guide to decisions about what to do.

But now suppose there are other good moral rules, such as ones that prohibit certain kinds of act, ones that permit some degree of preferential concern for oneself, and ones that require some degree of preference for one’s friends and family in one’s decisions about how to allocate one’s time, attention, and other resources. If there are these other good rules, then at least sometimes impartially applying and complying with them will conflict with letting impartial benevolence dictate what to do. More importantly, we can reject impartial benevolence as *the* direct guide to decisions about what to do while endorsing impartial application of good (first-order) moral rules.

Likewise, rejecting impartial benevolence as *the* direct guide to decisions about what to do does not entail rejecting impartial assessment of (first-order) moral rules.

Section 1 of this chapter argues that impartiality in the application of good moral rules is always appropriate. Section 2 argues that impartial benevolence as a direct guide to decisions about what to do is appropriate only sometimes. Section 3 argues that impartiality in the assessment of rules is or is not appropriate—depending on how plausible the impartially selected rules are.

I. Impartial Application of Rules

Many people closely associate morality and impartiality. I think a large part of the reason for this is that the impartial application of good moral rules is virtually always appropriate. If some rules really are good ones, and if these rules are being applied, then they should be applied impartially.

I am not saying that the impartial application of *bad* rules is appropriate. Indeed, with respect to *terrible* rules, *no* application of them is appropriate. And their impartial application might be, on balance, worse than their partial application.¹

Of course, there can be enormous uncertainty and debate about which possible moral rules are good ones, and about what makes them good. There are different plausible views about these things, especially about what makes rules good. I shall come back to these issues later when I discuss impartial assessment of (first-order) moral rules. At the moment, however, I want merely to make the point that the impartial application of good moral rules, whatever such rules turn out to be, is virtually always appropriate.

I should also acknowledge that some philosophers might reject the idea that there are good moral rules because they think that *rules* cannot help but be too coarse-grained to be good. I hope that such philosophers have something instead of rules—e.g. defeasible generalizations, default reasons, or hedged principles—the substitution of which will allow to go through much of what I argue below.

Now what does impartially applying a rule involve? Bernard Gert puts forward an account of *impartial treatment*:

*A is impartial in respect R with regard to group G if and only if A's actions in respect R are not influenced at all by which member(s) of G are benefited or harmed by these actions.*²

I take it that A's treatment of members of group G must be either random or patterned. If it is intentionally patterned, then it is probably guided by a rule. If it is guided by a rule, then when is the rule being impartially applied? Gert's account suggests an answer: a rule is being impartially applied with regard to a group if and only if its application is not influenced by which members of the group are benefited or harmed.

¹ This is a point Joel Feinberg makes: see his *Social Philosophy* (Englewood Cliffs, NJ: Prentice-Hall, 1973), pp. 106–7.

² Bernard Gert, *Morality* (New York: Oxford University Press, 1998), p. 132. I have discussed Gert's view of impartiality before. See my *Ideal Code, Real World: A Rule-consequentialist Theory of Morality* (Oxford: Oxford University Press, 2000), pp. 23–4. There I criticized Gert for not building consistency into impartiality. I advance different criticisms in this chapter.

With respect to many rules, this account of impartial application is attractive. The rules I have in mind are ones that, because of vagueness or complexity, do not always have clear implications. In effect, applying such rules requires lots of judgement, and leaves some degree of ‘wobble room’. An example might be the rule ‘The product of a collective enterprise should be divided in proportion to people’s contributions to the enterprise’. There is room for reasonable dispute about what makes one contribution greater than another.

To illustrate, suppose that Ivan applies the rule ‘reward people in proportion to their contribution’, and his allocations of reward are within the band of reasonable interpretations of ‘contribution’. But suppose his interpretation ranks contributions of strength over contributions of planning and innovation. And suppose he does this because he likes the strong more than the intelligent. In this case, when Ivan applies the rule ‘reward people in proportion to their contribution’, he is not doing so impartially.

Gert’s own list of moral rules comprises prohibitions on how we treat others.³ His first five rules prohibit us to cause any of five kinds of harm to others: death, pain, disability, loss of freedom, and loss of pleasure. His other five rules prohibit deceiving, cheating, breaking our promises, disobeying the law, and failing to do our conventionally determined duty. He insists that to comply with these prohibitions is to act impartially with respect to these rules. For example, if you never break your promises to anyone, then you have impartially complied with the rule against breaking promises. I accept that this is an example of impartially applying a moral rule.

What I cannot accept, however, is that impartially applying rules must involve not being influenced by which people are benefited or harmed. Consider the rule:

When you could devote your own time, attention, or other resources either to benefiting your friends and family or to benefiting people to whom you have no special connection, and when the benefit given would be about the same size, you should choose to benefit your friends or family.

In other words,

Given that you are selecting from the group of people who would benefit a fixed amount if you devoted your own time, attention, or other resources to them, you should choose your friends or family as the beneficiaries.

Rules such as this tell you to be influenced by which members of the relevant group will benefit from your action. There are many such rules. Another

³ Gert, *Morality*, chs 5, 7, 8.

example is one mentioned earlier: 'The product of a collective enterprise should be divided in proportion to people's contributions to the enterprise'. Such rules can be applied impartially.

Gert seems to have two possible responses to this line of thought. One possible response is to stick with the idea that to treat members of a group impartially is to treat them according to a rule that takes as irrelevant who benefits and who is harmed. If we stick with that idea, then we must say that you cannot treat members of a group impartially if you are following a rule specifying which members of the group qualify for benefits or harms, or for *more* benefit or harm. If you donate your own time and energy to helping your friends rather than to helping strangers, you have not treated your friends and the strangers impartially. But you may have impartially applied the rule 'donate your own time and energy to helping your friends rather than to helping strangers'.

Gert's other possible response is to give up the idea that to treat members of a group impartially is necessarily to treat them without regard to who benefits and who is harmed. If we give up that idea, then we can say that treating members of a group impartially is not incompatible with following a rule specifying which members of the group qualify for benefits or harms, or for *more* benefit or harm.

I favour allowing that impartial treatment of a group *can* be compatible with treating members of that group in accordance with a rule specifying which members of the group are so qualified. For example, acting impartially towards the company's workers with respect to rewarding productivity requires making decisions that give greater benefits to those of the company's workers who have been more productive. What impartial treatment of the company's workers requires is being guided by their past productivity rather than by other things, such as how much you like them.

How then should we characterize the impartial application of rules? Impartial application of a rule consists in being guided solely by the distinctions identified as relevant by the rule. Some rules make the issue of who is benefited or harmed irrelevant. This is true of rules against lying, stealing, and breaking the law. But other rules distinguish between others by picking out who qualifies for benefits or harms, or for *more* benefit or harm. Examples are rules about benefiting friends and family members over others when one is allocating one's own resources and the size of the benefits (by 'benefits' I mean positive contributions to welfare or personal good, not merely material goods) would be the same, and rules about rewarding productivity.

I turn now to a different aspect of Gert's account of impartial treatment of rules. Gert holds that someone can be morally impartial either in obeying rules

or in violating them. He thinks there are cases where violation of a rule is justified, e.g. lying to the murderer at the door. And he offers us a criterion for determining whether the rule was violated, either, on the one hand, impartially or, on the other hand, in order to make an exception for oneself or someone else. His view is that moral impartiality is satisfied as long as whoever violates a moral rule is willing for everyone to know that this kind of action is ‘publicly allowed’.⁴ Gert’s suggestion is that impartial violations of a moral rule entail the proposal of a *new, more specific* public moral rule indicating that in such circumstances violations of the older, less specific rule are permissible.

However, impartially applying and complying with a public rule might not yet be enough for moral impartiality. To be sure, the idea that the same rules apply publicly to everyone is associated with impartiality. However, advocating a single set of rules for public application to everyone hardly qualifies the advocate as impartial. Suppose Stephan wills and advocates public acceptance of a certain code of rules, but does so because he and his sub-group will be advantaged by public acceptance of this code. Stephan may apply these rules completely impartially. If he does, then he impartially applies rules that he sincerely wants to be publicly accepted. Still, to impartially apply rules that one sincerely wants to be publicly accepted is compatible with being very partial at the level of assessing rules.

Admittedly, publicly advocating rules that one publicly acknowledges to be particularly favourable to a group of which one is a member will often be pointless or even alienating. But publicly advocating rules that one publicly acknowledges to be particularly favourable to a group of which one is a member can be successful. Suppose you are hardworking and publicly advocate rules that are particularly favourable to the hardworking. These rules might well be accepted, and partly on the basis of your advocating them. Nevertheless, if you endorse these rules at least partly because you benefit from them, your endorsement of such rules is hardly unbiased. And other people’s endorsement of these rules might also spring from partiality toward the hardworking.

2. Impartial Benevolence as the Direct Guide to Decisions about What To Do

By impartial benevolence as a direct guide to decisions about what to do, I mean impartial benevolence as the direct and sole determiner of everyday

⁴ Gert, *Morality*, pp. 151–2, ch. 9.

practical decisions. By impartial benevolence, I mean an equal concern for the good of each. And by equal concern for the good of each, I mean treating a benefit or harm to any one individual as having the same moral importance as the same-size benefit or harm to any other individual.

What would it be to have impartial benevolence as the direct determiner of your everyday decisions? Benefits to anyone else would count in your reasoning for no less than the same-size benefits to you. Benefits to strangers would count in your reasoning for no less than the same-size benefits to your partner, child, or mother. So, if you recognized that donating most of your wealth to Oxfam would benefit the starving more than keeping it for yourself would benefit you, you would donate it. Indeed, you would go on giving your money, time, and effort to others as long as you thought the benefits others were getting were at least a little more than the benefits you were losing.

You might know more about how to benefit your family and friends than you know how to benefit strangers. Thus you might attend to your family and friends more than to others—but not because you have greater concern for your family and friends. On the contrary, whenever you were sure that doing something for a stranger would benefit the stranger at least a little more than doing the same thing for yourself or your family member or friend, you would benefit the stranger. If you could save three lives by giving to one person one of your kidneys, to the second person the other of your kidneys, and to the third person your heart, you would do so.

So far, I have been referring to what you do with your own money, time, effort, body parts, and other resources. But, if impartial benevolence really did determine *all* your everyday decisions, then presumably you would also be disposed to direct other people's resources in whatever way would maximize aggregate net benefit. If you could get away with channelling some of your employer's money to Oxfam, you would. If you could get away with channelling some of your friend's money to Oxfam, you would. If you could save two children by arranging for your child's kidney to go to one and her heart to the other, you would.

Absurd? Yes. Unfamiliar as an ethical ideal? Well, notoriously, we have William Godwin calling for the sacrifice of his mother to save the important do-gooder Archbishop Fénelon.⁵ And act-utilitarianism is routinely ridiculed for supposedly requiring agents to be prepared to do whatever it takes to maximize aggregate net benefit, impartially calculated.

⁵ Godwin, *An Enquiry concerning Political Justice, and its Influence on General Virtue and Happiness*, 2 vols (London: G. G. & J. Robinson, 1793).

W. D. Ross was right that act-utilitarianism ignores the ‘personal character of duty’.⁶ Ross was referring to the special relations each stands in to only some other people. The moral relations between friend and friend, or family member and family member, are, in many contexts, relevantly different from the moral relations between people with no special connection, and not just because we typically know more about our family members and friends than others. Even where we know equally well what would benefit a stranger and what would benefit a friend or family member, common moral opinion requires us to favour to some extent our own friends and family in the allocation of our own resources such as our time, energy, and material goods.

As John Cottingham has stressed, the duty of partiality pertains to the allocation of the agent’s *own* resources. There is no implication that the agent is permitted, much less required, to commandeer *someone else’s* resources for the benefit of the agent’s child. As John Cottingham puts it,

What *is* wrong with the Pope giving the Red Hat to his nephew, the judge deciding a case in favour of her cousin, the civil servant giving a contract to his pal, the admissions officer reserving a place for her friend’s daughter, is that such acts involve disrespect for the resources or rights of others. If I am working for the Church, or the Courts, or the Government, or the University, then the goods in question are not *mine*, to assign at will: I control the relevant good in trust for the institution that employs me, and I am no more justified in bestowing them on my favourites than I am justified in dishing out someone else’s cream to my cat, or ‘giving’ someone else’s bicycle to my child.⁷

In short, aiming to benefit oneself or one’s family, friends, or other associates is often off limits.

Indeed, *constrained* impartial benevolence as the determiner of practical decisions is absolutely mandatory in certain contexts. The person running CARE or the UN’s world food programme is charged to consider the welfare of everyone, considered impartially. For someone in such a role, the welfare of people of one religion, race, or region matter just as much as the people of any other religion, race, or region. And even in more restricted contexts, say a country or a county, various kinds of officials are charged to have the best interests of all their constituents equally at heart.

Nevertheless, *unconstrained* impartial benevolence would be inappropriate even in such roles. Even if you occupy such a role, you are *not* charged to do just *whatever* would maximize aggregate welfare. Some possible actions

⁶ W. D. Ross, *The Right and the Good* (Oxford: Clarendon Press, 1930), p. 19.

⁷ Cottingham, ‘The Ethical Credentials of Partiality’, *Proceedings of the Aristotelian Society* 98 (1998), pp. 1–21, at p. 11.

are ruled out by deontological prohibitions on, for example, murder, torture, robbery, fraud, etc. However, when you occupy one of various official roles and are choosing among possible actions none of which is ruled out by deontological prohibitions, then impartial benevolence should determine what you choose.

In contrast, where the resources you are allocating are your own, some degree of partiality is permitted or required. Yet, even here, the amount of permissible partiality is not *infinite*. You should give a benefit to your child or friend even when you could instead give a stranger a *somewhat* greater benefit. But, if you could either give to your mother an additional *minute* of happy life or give to someone with whom you have no connection an additional *decade* of happy life, you would be wrong to choose the tiny benefit for your mother rather than the very large benefit for the stranger.

Where exactly is the line dividing permissible from impermissible degrees of favouritism? A large grey area looms. On the one side, there are cases where the gap in the size of benefits is small enough to make it clear that you should favour those with whom you have special relations. On the other side, there are cases where the gap is big enough to make it clear that you shouldn't choose the smaller benefit for your friend or family member over the larger benefit for the other person. Between these two sets of cases is a large grey area, which itself has grey borders.

Complexities about patterns of decisions are also relevant. There are over a billion people in the world living in terrible poverty, with the usual concomitants—hunger, disease, low life expectancy, high infant mortality, etc. Imagine that, on every occasion when I could spend some time with my friends or instead devote that time to raising money for the worst off, I chose to raise money for the worst off. In doing this, I might on each occasion be helping to produce a vastly greater good than I would produce by using the occasion to spend time with my friends. But I would probably be starving my friendships of the sustenance they must have in order to survive. Surely morality wouldn't typically require of people a pattern of decisions that effectively deprives them of friendships.

Most humans have an immediate and intense special concern for themselves, their family, and their friends. Such partiality is widespread in the animal kingdom, and the theory of evolution easily explains this. Since humans evolved from animals, it is hardly surprising that a fairly high degree of partiality is instinctual in humans. Furthermore, as Cottingham comments, while such partiality might be flawed, it 'certainly is not self-evidently so'.⁸

⁸ Cottingham, 'The Ethics of Self-Concern', *Ethics* 101 (1991), pp. 798–817, at p. 814.

Another point of Cottingham's that I accept is that personal relationships, including family ones, are one of the central elements of the good life.⁹ He focuses on the sense of fulfilment that people get from their personal relationships. While hardly rejecting that sense of fulfilment, I think personal relationships also have non-hedonic value as an element of personal good. The value of personal relationships provides some grounds for concluding that, if impartial benevolence as a guide to everyday decisions is incompatible with personal relationships, so much the worse for impartial benevolence as a guide to everyday decisions.

Even if there were not such high value in personal relationships, natural human partiality might be so deep in our genes that there is little point in claiming that we ought to be impartial. And even if natural human partiality can be suppressed to the point of elimination, how much effort and energy and stress would be involved in this suppression? And how often would the effort and energy and stress recur? Absent genetic engineering, the costs of stamping out partiality in any generation would be high. Whatever the education and habituation of one generation, the next generation would come out of the womb pretty much the same as previous generations have—that is, with a very strong predisposition to partiality. So partiality would have to be stamped out in each new generation (unless it was eliminated by genetic engineering).

In this section, I noted (a) that common-sense morality endorses constrained partiality, (b) that, very plausibly, personal relationships, which are hugely valuable, wouldn't be possible without partiality, and (c) that partiality would anyway be costly to stamp out.

3. Impartial Assessment of (First-order) Moral Rules

As I indicated in the first section, almost everyone believes that morality requires impartiality in the application of good moral rules. And those who reject the idea that morality requires impartiality in the application of good moral rules reject this idea because they reject the idea of good moral rules, not because they reject the requirement to apply good moral rules impartially if there are any good moral rules.

We have also seen that impartiality construed as impartial benevolence is required if one is occupying certain official roles. And yet, again, impartial

⁹ 'Ethical Credentials of Partiality', p. 10; 'Ethics of Self-Concern', p. 816; 'Partiality and the Virtues', in R. Crisp (ed.), *How Should One Live?* (Oxford: Oxford University Press, 1996), pp. 57–76, at pp. 63–4.

benevolence as the determiner of *all* practical decisions is incompatible with people's intense special concern for themselves, their family, and their friends. Intense special concern for family and friends certainly seems permissible and even obligatory. Intense special concern for oneself certainly is natural and seems permissible, even if not obligatory. So, is impartial concern appropriate only when one occupies certain official roles?

Consider the two-level approach to impartial concern. What I mean by this is the idea that we might find that impartial assessment of (first-order) moral rules ends up favouring rules that give very considerable scope to partiality.

I contended earlier that impartial application of a rule consists in being guided solely by the distinctions and relations identified in the rule. And I pointed out that rules often direct that greater benefits go to some (e.g. the productive) than to others (e.g. the unproductive). But why are some rules the right ones? Why are the distinctions made by such and such rules the right distinctions to make? Here impartiality returns. The right distinctions, the right rules, are the ones that are impartially defensible.

What is it for a rule to be impartially defensible?

One idea is that for a rule to be impartially defensible is for it to be defensible from an agent-neutral point of view. If this idea is right, then your evaluating rules impartially is your evaluating them apart from any special attachments of yours. In your assessment of rules, you would not give extra weight to benefits that the rules produce for you, for your friends, for your family, etc.

In the assessment of rules, agent-neutrality does seem much more appealing than agent-relativity. The agent-relative assessment 'Everyone's accepting these rules is good because this maximizes benefits for me' is utterly unconvincing. So is 'everyone's accepting these rules is good because this maximizes benefits for my group'. Agent-neutral assessment effectively eliminates bias towards oneself, one's group, and indeed anyone with whom one has some special connection. The elimination of such bias at the level of fundamental assessment of possible moral rules certainly seems desirable.

One objection to the thesis that rules are impartially justified as long as they are selected by agent-neutral assessment is that there is very little prospect that the requirement that rules be evaluated agent-neutrally will by itself lead to convergence between us. Why is this an objection? One of the appealing prospects of impartial evaluation of rules is that, unlike agent-relative evaluation of rules, impartial evaluation of rules might, at least if we are all equally well apprised of the expected consequences, produce convergence among us. Admittedly, some philosophers (Gert, for example) are not optimistic about such convergence. They contend that, at least in many areas, impartial, rational, well-informed assessment cannot reasonably be

expected to produce it. In contrast, I hold out the hope that impartial, rational, well-informed assessment would produce convergence. Perhaps this is in part because of the central role I take moral principles to play in interpersonal justification and conflict resolution.

Let me illustrate how the requirement that rules be evaluated agent-neutrally might not lead to convergence between us. Suppose you think that the best songwriters are the most admirable people. And suppose that you evaluate rules primarily in terms of their effects on the best songwriters. Now your assessment of alternative possible moral rules is agent-neutral. It is *not* that, *because of some special connection to you*, you focus on how much the best songwriters benefit from this or that set of rules. For example, it is not that you yourself are a songwriter, or have a songwriter as a family member or friend. Nor is it the connection between songwriters and your sentiments that underwrites your focusing on the advantages and disadvantages to songwriters of this or that set of rules. Rather, the best songwriters just are the most admirable people, according to you. So that is how you evaluate rules.

Now suppose that the most admirable people, according to me, are sports stars. Again, suppose it is *not* that, *because of some special connection to me*, I focus on how much sports stars benefit from this or that set of rules. I'm not a sports star, nor is anyone with whom I have a special connection. So, in evaluating rules by their effects on sport stars, I am evaluating these rules agent-neutrally.

The group of people to whom you are giving priority in your assessment of rules overlaps very little if at all with the group of people to whom I'm giving priority in my assessment of rules. Hence, the set of rules ranked highest by you almost surely isn't the set ranked highest by me. In this example, you and I are both evaluating rules agent-neutrally, but without much serious chance of converging on the same set of rules.

Another and perhaps even more devastating objection to the thesis that rules are impartially justified as long as they are selected by agent-neutral assessment is that evaluating rules agent-neutrally is far from sufficient to obtain plausible results. Your agent-neutral assessment of rules privileges songwriters. For example, the rules you favour might place little demand on songwriters to do good for others, because songwriters need to be left as free as possible to do their songwriting. And my agent-neutral assessment of rules privileges sports stars. The rules I favour might place little demand on sports stars to do good for others, because they need to be left as free as possible to focus on their sports. But such concessions for songwriters or sports stars are intuitively wrong. So neither of us ends up with a plausible set of rules, though we both did our respective assessments agent-neutrally.

In short, evaluating *impartially in the sense of agent-neutrally* does *not* entail evaluating with *equal concern* for everyone.¹⁰ The examples of assessment biased towards the best songwriters and assessment biased towards sports stars are merely special cases of agent-neutral perfectionism. In the spirit of offering a less idiosyncratic agent-neutral perfectionism than the one privileging songwriters and than the one privileging sports stars, someone might hold that rules are to be evaluated in terms of the effects on the most talented, whether they be in the arts, literature, science, engineering, business, sports, military activities, etc. This evaluative stance gives priority to benefits for the talented. It is a kind of elitism of the talented.

As I said a moment ago, agent-neutral assessment of rules is better than agent-relative assessment of them. The agent-neutral assessment ‘Whether or not I or anyone connected with me is talented, everyone’s accepting these rules would be good because this would benefit the most talented’ is better than the agent-relative assessment ‘Everyone’s accepting these rules would be good because this would maximize benefits for me, or for my group’.

Yet elitist assessment of the basic rules of morality is unacceptable. We’ve seen that elitist assessment can be agent-neutral. So an evaluative stance’s being agent-neutral isn’t enough to make it acceptable. Thus, if ‘impartial’ is taken to mean ‘agent-neutral’, we have to reach the conclusion that an evaluative stance’s being impartial isn’t enough to make the evaluative stance acceptable. Far better if we can find a meaning for ‘impartial’ that does make an evaluative stance’s being impartial enough to make it acceptable.

And, clearly, there is a perfectly obvious and straightforward sense in which the elitist evaluative stance isn’t impartial. This elitist assessment gives priority to benefits to the most talented (or most accomplished). It does not count a benefit to the untalented for as much as the same-size benefit to the talented. In this obvious sense, elitism is not impartial as between the talented and the untalented.

In fact, the most obvious form of impartial concern is utilitarian concern, i.e. concern that takes a benefit to any one individual to have exactly the same importance as the same size benefit to any other individual.

One worry about this sort of impartial assessment of rules is that there will be slippage between the levels. At the deepest level, rules are to be assessed impartially. The set of rules with the highest expected net benefit will require certain kinds of partiality in everyday decision-making about how to allocate one’s own time, energy, money, etc. The set of rules with the highest expected

¹⁰ This point is another that Cottingham has noted; see his ‘Ethical Credentials of Partiality’, pp. 4–5.

net benefit will also include a rule allowing agents to focus their lives on their own projects. But the worry is that the impartialism in the assessment of rules will seep into everyday decision-making.¹¹

This worry is misplaced. How many people do you know who manifest too much impartial concern in their day-to-day decision-making? Selfishness, strictly confined generosity, concern for the neighbour that doesn't extend to the people on the far side of the neighbour, concern for members of one culture that doesn't extend to members of other cultures—all these are very common. In contrast, people who give their own welfare and the welfare of their near and dear too little consideration in comparison with the welfare of strangers can hardly be said to be teeming in the streets. That is hardly surprising. After all, natural inclination and the desire for gratitude from near and dear line up on the side of partiality. Impartial concern is the side having to fight an uphill battle.

Another objection to impartial *utilitarian* assessment of rules is that the best rules are instead the ones that maximize expected value where this is calculated by giving some degree of *priority to benefits to the worse off*. This view about what selects the best rules is called weighted prioritarianism.¹²

How much priority do benefits to the worse off get in weighted prioritarianism? Well, definitely *less* than *absolute* priority. But how much less? Different weighted prioritarians will give different answers. And maybe no weighted prioritarian will be very certain or specific about their answers. There is thus an understandable tendency for weighted prioritarianism to be somewhat vague about how much priority should be given to benefits for the worse off.

Admittedly, at least some such versions of weighted prioritarianism generate a very plausible ranking of possible sets of moral rules. But perhaps the intuitively plausible implications that these versions of weighted prioritarianism has are the flip side of their vagueness. (Or, at least, the vaguer a principle is, the harder to find cases where the principle's implications clearly are intuitively implausible.)

If weighted prioritarianism is not vague, it specifies a specific degree of priority. But any particular specific degree of priority will be difficult to defend as better than every other degree.

¹¹ For one example of this objection, see Cottingham, 'Morality, Virtues and Consequences', in D. Oderberg and L. Laing (eds), *Human Lives* (London: Macmillan, 1997) pp. 128–43, at p. 139.

¹² For influential discussions, see Joseph Raz, *The Morality of Freedom* (Oxford: Clarendon Press, 1986), p. 227; Thomas Nagel, *Equality and Partiality* (New York: Oxford University Press, 1991), ch. 7; Derek Parfit, 'Equality and Priority', *Ratio* 10 (1997), pp. 202–21; and David Miller, *Principles of Social Justice* (Cambridge, MA: Harvard University Press, 1999), pp. 223–5. For an important recent discussion, see Marc Fleurbaey, Bertil Tungodden, and Peter Vallentyne, 'On the Possibility of Nonaggregative Priority for the Worst Off', *Social Philosophy and Policy* 26 (2009), pp. 258–85.

Another argument against weighted prioritarianism is that it cannot serve as an *impartial* foundational principle for morality. Though agent-neutral, weighted prioritarianism gives priority to benefits for the worse off, and in this way is a form of partiality towards the worse off, much as forms of elitism are forms of partiality towards an elite. If elitism is not a form of impartiality, weighted prioritarianism is not either.¹³

The argument I have just rehearsed takes impartial benevolence—in which benefits or harms to any one individual have the same importance as the same size benefits or harms to any other individual—as the most natural form of impartiality in the assessment of rules. But that view can be challenged. While agreeing that rules should be assessed impartially, we might think that this requires them to be assessed in terms of justice or fairness as well as, or even instead of, in terms of net aggregate benefit. There is a variety of views about what constitutes justice or fairness. One prominent view is that justice or fairness requires that each gets what he or she deserves. Another is the view that each gets what he or she needs. Another is the view that the worst-off position should be as good as possible. Weighted prioritarianism is another rival here. And so on.

The point in mentioning this variety of views is that each of them makes a distinction that itself needs defence. What makes someone morally deserving or undeserving? What are needs and why is satisfying them, rather than (say) increasing net aggregate benefit, pivotal? Why is a benefit for someone who is worse off better than the same-size benefit for someone who is better off?

Most theorists answer this question by arguing that the distinction in question is a (or the) focus of justice or fairness. Debate then ensues about the plausibility of such claims. I now accept that no appeal to impartiality by itself can resolve the question of which agent-neutral considerations and distinctions do, and which do not, have foundational importance. I hence withdraw my previous complaint against weighted prioritarianism that it fails as a form of impartiality.

The last objection to my approach I consider here is one of John Cottingham's. He complains that consequentialists 'see the value of individual lives as essentially derivative from their contribution to impersonally defined goodness'.¹⁴ He writes, 'Human lives are valuable not in virtue of how far they

¹³ This is an argument I first presented in *Ideal Code, Real World*, pp. 60–2. I lay out the argument at greater length in 'Up and Down with Aggregation', *Social Philosophy and Policy* 26 (2009), pp. 126–47, at pp. 141–2. In the present essay, however, I shall go on to take a different line on the relationship between impartiality and prioritarianism. And I am especially grateful to Brian Feltham, Peter Vallentyne, Andrew Williams, and Michael Otsuka for pushing me to this line.

¹⁴ 'Ethical Credentials of Partiality', p. 8.

contribute, individually or collectively, to some giant amalgam called “the good”, but in so far as they are lived in ways which make the short journey each of us has to undergo meaningful and precious.¹⁵

I don’t understand this objection. Impartial consequentialists typically take aggregate goodness to be made up of the welfare of individuals, and the welfare of individuals is determined by how well their lives go. Since the aggregate good is a function of the good of individuals, the good of individuals is prior to (rather than derived from) the aggregate good. If anything, the aggregate good is derived from the good of individuals.

My discussion of impartiality in the assessment of moral rules has concentrated on impartial consequentialist assessment of rules. But I do not mean to suggest that consequentialist impartiality is the only kind. Contractualism and Kantianism are often offered as accounts of fundamental moral impartiality. Most famously, John Rawls offered his original position and veil of ignorance as an alternative to utilitarian impartiality.¹⁶

As far as I know, the most plausible development of Kantian and contractualist lines of thinking appears in Derek Parfit’s recent work.¹⁷ Parfit argues that the most plausible form of Kantian contractualism holds that an act is wrong if and only if it is forbidden by principles that everyone can rationally will that everyone accept. Parfit goes on to argue that the only principles that everyone can rationally will that everyone accept are the ones whose universal acceptance would make things go best. If Parfit’s arguments are sound, Kantian contractualism leads to the same principles, or rules, that rule-consequentialism endorses. There is, of course, room for dispute, e.g. about which interpretation of the phrase ‘universal acceptance’ makes the theory come out most plausible. But suppose those disputes can be satisfactorily resolved without undermining Parfit’s arguments. In that case, the leading forms of impartial assessment of rules, i.e. Kantian/contractualist and rule-consequentialist forms of impartial assessment, will have been shown to converge on the same set of rules. That in itself would be a discovery of immense importance.

Suppose Parfit is correct that the leading forms of impartiality converge on a certain set of rules. It might be that the rules converged upon are seriously counterintuitive. Or it might be that these rules are instead intuitively plausible.

Let us consider those two possibilities. If the rules that the leading forms of impartiality converge upon are seriously counterintuitive, those who have

¹⁵ Cottingham, ‘Morality, Virtues and Consequences’, p. 139.

¹⁶ John Rawls, *A Theory of Justice* (Cambridge, MA: Harvard University Press, 1971), sections 3, 4, 5, 24, 30.

¹⁷ Derek Parfit, *On What Matters* (Oxford: Oxford University Press, forthcoming).

been rejecting impartial assessment of rules will feel vindicated. On the other hand, if the rules that the leading forms of impartiality converge upon *are* intuitively plausible, then this not only enhances the credibility of those rules but also undercuts the only credible objection to impartial assessment of rules, i.e. that such assessment leads to implausible rules.

That is, the following meta-ethical principle about evaluating moral theories seems overwhelmingly plausible:

For any two moral theories, if they are roughly equally good at cohering with independently credible intuitions about which possible rules are good ones, and if one of these theories identifies a fundamental moral principle that provides impartial justification for these rules and the other theory doesn't, then the theory that identifies a fundamental principle that provides impartial justification for the right rules is better.¹⁸

¹⁸ I am grateful to the British Academy of Humanities and Social Sciences for a Research Readership, during which I worked on a project of which this is a part. For helpful comments on this chapter, I'm grateful to John Cottingham, Jonathan Dancy, David Estlund, Brian Feltham, John Kekes, Brian McElwee, William O'Brian, Andrew Williams, Jo Wolff, and two anonymous reviewers for Oxford University Press.