https://doi.org/10.1017/S0265052509090062 Published online by Cambridge University Press

UP AND DOWN WITH AGGREGATION*

By Brad Hooker

I. Challenges to the Very Idea of Aggregate Social Good

The most natural understanding of the expression "aggregate social good" is that it refers to the sum of the good of individuals within some social group. (Which social group might be obvious in the context of discussion, or might be contentious.) The most natural understanding of the good of individuals must be the welfare, or well-being, of individuals.

The only rival understanding to this welfarist one that I can think of is that the good of the individual is just the fulfillment of the individual's desires, and then the aggregate social good is just the aggregate fulfillment of the desires of the individuals in the group. Call this the *desire-fulfillment theory of individual and aggregate good.*

A confusing complication is that a familiar family of theories of individual welfare and of aggregate welfare holds that these consist in desire-fulfillment and aggregate desire-fulfillment. On the one hand, then, there are theories that do not mention welfare but nevertheless hold that individual and aggregate good consist in desire-fulfillment. On the other hand, there are theories that hold that good consists in welfare and that welfare consists in desire-fulfillment.

Whichever of these sets of theories is in play, there might seem to be a crippling problem for the very idea of aggregate social good. Is an aggregation of fulfillment of different people's desires intelligible?

To keep things simple, let us consider a two-person society. Suppose the question is whether there is more aggregate social good in an outcome where Jack gets his strongest desire fulfilled or in an outcome where Jill gets her strongest desire fulfilled. In each of these two outcomes, one and only one person has his or her strongest desire fulfilled. If an individual's good is measured only ordinally, then the two outcomes here must be equally good.

* This paper was one of the ones written while I held a research readership funded by the British Academy for the Humanities and Social Sciences. I am very grateful to the British Academy for the support. This paper was presented at the University of Reading and the University of Stirling. I am grateful to Peter Boltuc, Bryan Caplan, Andrew I. Cohen, Rowan Cruft, Jonathan Dancy, David Faraci, Brian Feltham, Guy Fletcher, Ian Folland, Russell Hardin, Iwao Hirose, Frances Kamm, Ambrose Lee, Mari Mikkola, Alan Millar, Elijah Millgram, Peter Milne, Jan Narveson, Alastair Norcross, Douglas Portmore, Anne Raustol, Jonathan Riley, David Sosa, Jussi Suikkanen, Peter Sullivan, Fiona Woollard, and Peter Vallentyne for helpful comments. I am especially grateful to Ellen Frankel Paul for pages of careful criticisms of an earlier draft.

doi:10.1017/S0265052509090062

© 2008 Social Philosophy & Policy Foundation. Printed in the USA.

https://doi.org/10.1017/S0265052509090062 Published online by Cambridge University Press

Intuitively, however, merely ordinal measures of individual good are inadequate. Everyday life repeatedly throws up examples where one person's strongest desire is but a weak one and someone else's strongest desire is very strong. Suppose Jack's strongest desire is fairly weak compared to Jill's strongest desire. A merely ordinal measurement of good would equate the outcome in which Jack gets his strongest desire fulfilled and the outcome in which Jill gets her strongest desire fulfilled. In such cases, merely ordinal measurements are manifestly mistaken, because merely ordinal measurement ranks these two outcomes equally. Cases such as this one cry out for cardinal measurement.

But cardinal measurements are highly problematic on any kind of desire-fulfillment theory of goodness. There are different ways of making cardinal measurements, and which way cardinal measurements are made can matter. On one way of making cardinal measurements, Jack might benefit more from the first possible outcome than Jill does from the second possible outcome. On another way of making cardinal measurements, however, Jill might benefit more from the second outcome than Jack does from the first outcome. Unless there is some nonarbitrary way of deciding which cardinal measurement to use, who benefits more in desire-fulfillment terms is underdetermined. This point stymied philosophers, economists, and policy analysts for many years.

But why was desire-fulfillment taken to be constitutive of welfare or good and therefore an (if not *the*) appropriate goal of moral and political decisions? I speculate that part of the answer is that desire-fulfillment seemed naturalistically verifiable and measurable (at least ordinally). Another part of the answer is that taking desire-fulfillment as a moral and political goal allows one to remain neutral about people's preferences and in this way respects autonomy. Let me explain.

In pointing to the supposed naturalism of taking desire-fulfillment as a goal, I mean that whether an act or state of affairs fulfills a desire is an empirical matter.[1] The most naturalistic sort of desire-fulfillment theory would be one that took good or welfare to be constituted by the fulfillment of actually existing desires. Those who are attracted by the supposed naturalism of taking desire-fulfillment as a goal are, I think, at least in part attracted by the following line of thought:

Premise 1: What desires a person actually has is a matter of discoverable empirical fact.

[1] This is not the place to get bogged down in debates about what naturalism is. Although the matter is controversial, I will assume here that natural properties are ones that can be known by empirical means and thus are appropriate subjects for investigation by the natural and social sciences. See G. E. Moore, *Principia Ethica*, 2d ed. (Cambridge: Cambridge University Press, 1993), 13; David Copp, "Why Naturalism?" *Ethical Theory and Moral Practice* 6, no. 2 (2003): 179–200, at p. 181; cf. Nicholas Sturgeon, "Moore on Ethical Naturalism," *Ethics* 113, no. 3 (2003): 528–56, at pp. 534–38.

Premise 2: Whether a state of affairs would fulfill a person's desires is also a matter of discoverable empirical fact.
Premise 3: A person's good is constituted by the fulfillment of his or her actual desires.
Therefore: A person's good depends only on discoverable empirical facts.

But there are problems with every premise of this argument.

The first premise provokes three lines of objection. First, what desires a person has might be so fluid or indeterminate that they are not in practice discoverable. Second, people sometimes do not know what they desire. Third, even if they do know, they might try very hard to suppress or at least camouflage their desires. Thus, even if it is possible *in principle* to discover what desires people have, it may not be possible *in practice* to discover what desires they have.

Now turn to the argument's second premise. Consider the question of whether a state of affairs would fulfill someone's actual desire *when that desire employs an evaluative concept*. Suppose someone desires that her life contain a truly significant achievement. The object of her desire is not to achieve merely something that she *thinks* is significant. What she thinks significant might be truly significant, but equally it might not be. What the person in our example really wants is to achieve something truly significant. Now how can it be a purely discoverable empirical fact whether her life does contain a truly significant achievement? The implication of such examples is that the question of whether a state of affairs would fulfill someone's actual desire where that desire itself employs an evaluative concept is not a purely empirical question.[2]

Let me move on to the point (related to the third premise) that what people *actually* desire (even if this is determinate and discoverable) cannot be the focus of an adequate theory of welfare. Actual desires might depend on false empirical beliefs, faulty logic, failures of imagination, etc. In these ways at least, actual desires might be based on mistakes. So actual desires cannot be the ground of welfare or good. An adequate theory of welfare or good *must* idealize the ground to at least some extent.

Such idealization will at least correct false empirical beliefs, faulty logic, and failures of imagination. So someone's "informed" desires are the ones that he or she would have after gaining all relevant empirical information, after any faulty logic is corrected, and after adequately imagining

[2] For a related argument that a desire-fulfillment theory of good reasons for action does not have the naturalistic purity some philosophers have supposed it has, see Jean Hampton, *The Authority of Reason* (New York: Cambridge University Press, 1998), esp. chap. 5.

https://doi.org/10.1017/S0265052509090062 Published online by Cambridge University Press

https://doi.org/10.1017/S0265052509090062 Published online by Cambridge University Press

the relevant scenarios. And such "informed" desires take precedence over actual ones.

To illustrate, suppose Jack actually desires to marry Jill, in part because he believes she finds him attractive. But that belief is false. If he knew the truth that she finds him unattractive, he would not desire to marry her. This is precisely the sort of case that makes the informed-desire theory seem superior to the actual-desire theory.

Now, what kind of information can come into play here? Those who advocate an informed-desire theory appeal to relevant empirical information but not evaluative information. In this way, the informed-desire theory retains a naturalistic, empiricist cast. And I think this element of the theory has attracted many philosophers, economists, and policy analysts to it.

Another element is what I have called the neutrality of desire-fulfillment theory. What I mean is that desire-fulfillment theory passes no judgment on the content of desires—as long as those desires are not based on false empirical beliefs, faulty logic, failures of imagination, and the like. In this way, desire-fulfillment theory respects what John Harsanyi called "preference autonomy."[3]

Desire-fulfillment theory may have been dominant for many decades, but the arguments in its favor are unconvincing. As I explained above, the question of whether a state of affairs would fulfill someone's actual desire is perhaps not a purely empirical question when that desire employs an evaluative concept. Even if desire-fulfillment theory gives a purely naturalistic account of welfare or good, perhaps other theories, including objectivist ones, can also be equally naturalistic.[4] Furthermore, even if desire-fulfillment theory is not only entirely but also uniquely naturalistic, whether naturalism about evaluative concepts (such as the concepts of good and of welfare) is correct is highly controversial.[5]

Though I accept that preference autonomy is an attractive feature of desire-fulfillment theory, I do think it is not attractive enough to outweigh the problems with desire-fulfillment theory. Everyday life provides us with examples of bizarre actual desires—including widespread ones. But even empirically informed, logically reasoned desires *could* be bizarre. Standard examples from the literature are the person who desires a saucer of mud for its own sake, the person who desires to count all the blades of

[3] John Harsanyi, "Morality and the Theory of Rational Behaviour," in Amartya Sen and Bernard Williams, eds., *Utilitarianism and Beyond* (Cambridge: Cambridge University Press, 1982), 39–62, at pp. 60–61.

[4] I explain this in my essay "Theories of Welfare, Theories of Good Reasons for Action, and Ontological Naturalism," *Philosophical Papers* 20, no. 1 (1991): 25–36.

[5] Recent defenses of nonnaturalism include T. M. Scanlon, *What We Owe to Each Other* (Cambridge, MA: Harvard University Press, 1998), 58–61; Russ Shafer-Landau, *Moral Realism* (New York: Oxford University Press, 2003); and Derek Parfit, *Climbing the Mountain* (Oxford: Clarendon Press, forthcoming).

grass in the lawns along a street, and the person who desires to turn on as many radios as possible.[6]

My point is that the fulfillment of such desires does not, as such, add to a person's welfare. If the person would get pleasure from the fulfillment of such desires, this pleasure is a benefit. If the person would feel frustrated or distressed from not fulfilling such desires, the felt frustration or distress would be a cost. But the fulfillment of desires for ends that are not desirable constitutes no addition to welfare or good.[7]

Those last remarks are not meant to suggest that some sort of hedonism is the right account of welfare or good. In fact, hedonism is too monistic a theory of welfare and good. Friendship, achievement, and important knowledge, for example, are of course typically valuable as means to pleasure and the avoidance of pain—but they also have value beyond whatever their hedonistic payoffs may be. For this reason, I accept an objective-list theory of welfare.[8] The list contains at least friendship, achievement, important knowledge, and pleasure.

I have not argued here for the objective-list theory over hedonism. Rather, I have briefly reviewed the arguments for and against desire-fulfillment theory, coming to the conclusion that there is too much going against desire-fulfillment theory for it to be acceptable. The relevance of this conclusion is that the main qualm about the very idea of aggregate social good comes from combining the idea that the best account of individual good is some sort of desire-fulfillment theory with the idea that the choice of a metric for cardinal measurements of individual good is arbitrary. Once the desire-fulfillment theory is jettisoned, the main qualm about the very idea of aggregate social good goes too.

But then how is aggregate social good to be calculated on a hedonistic or objective-list theory of individual good? The answer, I concede, is that very often judgments about these matters are speculative, highly contest-

[6] See Elizabeth Anscombe, *Intention* (Oxford: Blackwell Publishing, 1957), 70; John Rawls, *A Theory of Justice* (Cambridge, MA: Harvard University Press, 1971), 432; Warren Quinn, *Morality and Action* (New York: Cambridge University Press, 1993), 236; Roger Crisp, *Mill on Utilitarianism* (London: Routledge, 1997), 56; and Robert Audi, *Ethical Theory and Moral Character* (New York: Oxford University Press, 1997), 77, 78.

[7] But what about desires for ends such as hitting small white balls into distant holes, as in golf, or, more generally, desires for sporting achievements that are in some sense the achievement of contrived goals? For a discussion, see Thomas Hurka, "Games and the Good," *Aristotelian Society Supplementary Proceedings* 80 (2006): 217–35; John Tasioulas, "Games and the Good," *Aristotelian Society Supplementary Proceedings* 80 (2006): 237–64; and my "The Meaning of Life: Subjectivism, Objectivism, and Divine Support," in Nafsika Athanassoulis and Samantha Vice, eds., *Morality and the Good Life: Essays in Honour of John Cottingham* (London: Palgrave, 2008): 184–200.

[8] See Derek Parfit, *Reasons and Persons* (Oxford: Clarendon Press, 1984), appendix I; John Finnis, *Natural Law and Natural Rights* (Oxford: Clarendon Press, 1980); Finnis, *Fundamentals of Ethics* (New York: Oxford University Press, 1983); David Brink, *Moral Realism and the Foundations of Ethics* (Cambridge: Cambridge University Press, 1989), 221–36; Crisp, *Mill on Utilitarianism*, chap. 3; Bernard Gert, *Morality* (New York: Oxford University Press, 1988), 92–94; and Richard Arneson, "Human Flourishing Versus Desire Satisfaction," *Social Philosophy and Policy* 16, no. 1 (1999): 113–42.

https://doi.org/10.1017/S0265052509090062 Published online by Cambridge University Press

able, and tremendously fallible. Furthermore, perhaps every judgment about individual or aggregate good is inherently to some degree imprecise.

However, judgments about individual and aggregate good can *sometimes* rightly be made with great confidence. Suppose we are asked about the aggregate social good of two different groups with an equal number of people. In the first group, life spans are short, disease is widespread, pain is unrelieved, and the struggle for short-term survival is all-consuming. In the second group, people are able to pursue long-term goals and relationships, have the facilities to prevent or relieve pain and disease, and have the luxury of studying questions other than how to protect their group. Given these facts about the two groups, we can be confident that the aggregate social good in the second group is considerably greater than that in the first. Of course, that is but an abstractly described thought-experiment. Nonetheless, in many specific real-life cases, those who are in a position to make judgments are rightly confident that the aggregate social good of one group is greater than that of another, or that the aggregate social good of a group under one set of circumstances is greater than it would be under a different set of circumstances.

Again, I concede that, very often, estimates of aggregate social good are impressionistic, speculative, highly contestable, and tremendously fallible. But that is no reason for ignoring comparisons of aggregate social good in the cases where confident (though imprecise) judgments can be made (and there are at least many such cases). There may well be *other* reasons for ignoring comparisons of aggregate social good in the cases where confident judgments can be made. I will consider those reasons in Sections III through V of this essay. Still, the point that needs to be stressed first is this: the fact that estimates of aggregate social good are often too speculative, and always to some extent imprecise, is not a good reason to dismiss them when they are well grounded.

II. The Case for Aggregate Social Good

Imagine a world with two billion people but also with a stomach cancer that kills everyone when he or she turns fifty (unless the person has died earlier of other causes). Call this *death-at-fifty stomach cancer.* One day, some forty-nine-year-old people happen to breathe in the vapors of a boiling fern. Their stomach cancer is completely cured. In addition, these people suffer no harmful side effects. Moreover, this fern is very easy to grow anywhere. So the survivors see that what saved them might save many others.

The survivors tell their doctors. The doctors do lots of tests on them. Encouraged by these tests, the doctors try out the fern vapor on a thousand other forty-nine year olds. Every single one is cured without any

https://doi.org/10.1017/S0265052509090062 Published online by Cambridge University Press

harmful side effects. Perhaps now, or after more tests, the doctors conclude that the fern-vapor treatment is likely to cure every forty-nine year old of the death-at-fifty stomach cancer. Should the doctors publicize this information?

Of course they should! But what precisely is the reason they should? Is not the reason that everyone gains?

If everyone gains overall, then it might be hard to see any reason against publicizing the information. Actually, in Section III, I will discuss some possible reasons for opposing the introduction of this cure even if everyone gains from its introduction. But I think that, even if there are such reasons for opposition, they are outweighed by the benefits to everyone of introducing the cure.[9] If at least one individual would be better off in outcome A than in outcome B, and if no one would be worse off in outcome B than in outcome A, then outcome A is Pareto-superior to B. Pareto-superior outcomes are very hard to resist.

But now suppose that dissemination of the fern-vapor cure is not beneficial to everyone. Perhaps some people were in businesses that made sense when everyone died at fifty but make no sense when life expectancy is extended to eighty. Perhaps the economic effects on some of these people are so bad that they now starve to death before they reach the age of fifty. So some people might lose from the introduction of the cure. Where this is the case, introducing the cure is not Pareto-superior.

What if those who would lose are only a few and would lose only a little per person, and those who would gain are many and would gain a lot per person? For the sake of argument, imagine that those who gain from the introduction of the fern-vapor treatment gain an average of thirty years of healthy life. And imagine that those who lose overall from its introduction lose an average of a few days of life. (To isolate relevant variables, let us assume that death-at-fifty stomach cancer is exactly as painful as starving to death and takes exactly as long.)

We can now imagine that the doctors face the following choice. If they introduce the fern-vapor cure, 1,999,900,000 people will get an average of thirty extra years of healthy life, but 100,000 will be on balance a little worse off. If they do not introduce the cure, the 100,000 will not lose out, but each of the 1,999,900,000 will miss out on the thirty years of increased life expectancy.

Someone might say that the reason the doctors should introduce the fern-vapor cure in this example is simply that this would benefit the greatest number of individuals. The suggestion, in other words, is that majoritarianism is what drives forward the decision in this case, not appeal to aggregate social good.

[9] For an influential discussion, see Joseph Raz, *The Morality of Freedom* (Oxford: Clarendon Press, 1986), 227.

https://doi.org/10.1017/S0265052509090062 Published online by Cambridge University Press

https://doi.org/10.1017/S0265052509090062 Published online by Cambridge University Press

Well, the case I described is one in which majoritarianism and aggregate social good favor the same decision. So the case is useless for deciding between majoritarianism and aggregate social good. However, we can think of cases where majoritarianism conflicts with aggregate social good.

Suppose now that, of the two billion people in this world, 999,999,500 would get an average of an extra thirty *years* of healthy life if the fern-vapor treatment is introduced, and 1,000,000,500 would end up dying a few *days* earlier than otherwise. Thus, in the example we are imagining now, a relatively narrow majority would be better off without the cure being publicized. But on the other side what is at stake for an almost equally large group is a very much greater gain per person.

Here I am focusing on *expected* net benefits. I admit that, in some cases, a few lost days for person A make a bigger difference to A's life than the extra thirty years for person B make to B's life. The few lost days for A might have been ones where important goals would have been finally reached, relationships consummated, and profound discoveries made. The extra thirty years for B might have been fairly empty ones. However, though there are such cases, they of course are historically rare and prospectively improbable. For any given pair of people, the probability is far greater that the benefit for one of gaining an extra thirty years of life will dwarf the cost for the other of losing a few days of life.

That sort of example, if persuasive, shows that majoritarianism is not correct. What is most beneficial to the majority might not be best all things considered.

One reaction to such cases as the one I have just described is that these are cases where aggregate social good, understood in an obviously utilitarian way, is really compelling. The expected aggregate net benefit of one alternative is far greater than that of the other alternative, despite the fact that this other alternative is better for the majority. And clearly there is some moral pressure in such cases to choose the alternative with greater expected aggregate net benefit.

Another reaction to such cases as the one I have just described is that what is actually doing the work is not the greater *aggregate* good in one alternative than in the other, but instead the greater good *per person.* Someone with thirty available years to lose has more to lose—typically at least—than someone with only a few days to lose. Thus, a pairwise comparison of individuals in the two possible outcomes suggests that the outcome in which the fern-vapor cure is disseminated is better.

But pairwise comparison is not always compelling. Consider the following case borrowed from Derek Parfit.[10] We have to choose between either:

[10] Derek Parfit, "Justifiability to Each Person," *Ratio* 16, no. 4 (2003): 368–90.

(1) Jack has fifty weeks of pain, and a thousand other people each have fifty-one weeks of pain, or
(2) Jack has no pain, and a thousand other people each have fifty-two weeks of pain.

If we are guided only by pairwise comparisons, we have to choose (2) rather than (1). The reason for this is that the difference to Jack's life made by the choice between (1) and (2) is fifty weeks of pain, whereas the difference to any other individual's life made by the choice is only one week of pain. But if we choose (2) over (1), we choose 52,000 weeks of pain over 51,050 weeks of pain. Furthermore, note that, in either case, each of the thousand is worse off than Jack. So if we do what is best for Jack, we are doing what is best for the person who is better off anyway. If we do what is best for Jack, we are promoting neither utility nor equality.

The reply to Parfit might be that, in our pairwise comparisons of the benefits and harms to different individuals, we should give far greater weight to the effects on the worst off. But now Parfit poses this choice:

(3) Jack has fifty weeks of pain, and Jill has fifty-one weeks of pain, or
(4) Jack has no pain, and Jill has fifty-two weeks of pain.

Were we to give far greater weight to the effects on the worst off, then in this case we would presumably have to choose (3) over (4), which would result in 101 weeks of pain rather than fifty-two weeks. This seems to be the wrong decision.

In sum, we intuitively believe that (1) is better than (2) and that (4) is better than (3). This conjunction of beliefs suggests that we believe, contra those who allow only pairwise comparisons, that aggregate benefits or harms can be morally decisive.

So far, I have defended the very idea of aggregate social good, construed it in broadly welfarist terms, admitted that our calculations of it are often impressionistic and fallible and always imprecise, and illustrated the kinds of cases in which aggregate social good, or aggregate expected net benefit, exerts powerful moral pressure on decision-makers. The next section considers objections to having aggregate social good be decisive.

III. Objections to Having Aggregate Social Good Be Decisive

Restrictions and rights often militate against doing what would maximize expected aggregate net benefit. There are moral restrictions on phys-

https://doi.org/10.1017/S0265052509090062 Published online by Cambridge University Press

ically harming others, and rights not to be physically harmed by others. There are also restrictions (albeit, of course, weaker ones) on stealing or damaging others' property, breaking promises, and telling lies. Again, there are corresponding rights, and these are correspondingly weaker than the right not to be physically harmed. Complying with any of these restrictions, respecting any of these rights, can get in the way of maximizing expected aggregate net benefit.

Notoriously, the same is true of many other rights. Perhaps this is most familiar in the case of the right to free expression, the right to vote, the right to legal representation, the right to face one's accusers, and the right to make binding contracts. In the case of each of these rights, their exercise very often fails to promote (let alone maximize) aggregate social good. Many people exercise their right to vote by voting for the wrong person. People frequently say things that cause more harm than good. They hire lawyers who succeed in producing outcomes better for their clients but worse in terms of aggregate social good. And binding contracts frequently have bad outcomes. Nevertheless, except perhaps in extreme circumstances, morality requires that these rights be respected even if respecting them in the individual case will predictably not serve aggregate social good.

Special obligations can also militate against maximizing aggregate social good. If you could use your own time, money, or other resources to benefit your friends or family, or instead to benefit others who have no special connection with you, then you should benefit your friends and family unless perhaps the amount others would benefit would be *much* more. For example, suppose you could use your own money to purchase an operation for your mother that would give her one extra year of healthy life. Or you could use the money to purchase an operation for a stranger that would give him two extra years of healthy life. Because of the special obligation you have to your mother, you should benefit her, even if aggregate social good would be a bit greater if you benefited the stranger instead.

Special obligations, others' rights, and restrictions on how we are to treat others—all these can put moral pressure on us in particular situations not to do what would maximize aggregate social good. Indeed, special obligations, rights, and restrictions can conflict with aggregation of more than one kind. I have been focusing on the potential conflict between maximizing aggregate social good and complying with restrictions, respecting rights, and fulfilling special obligations. But complying with restrictions, respecting rights, and fulfilling special obligations can also conflict with maximizing aggregate compliance with special obligations, rights, and restrictions.

This is because complying with restrictions, respecting rights, and fulfilling special obligations are *agent-relative* duties. A duty is agent-relative if it includes an essential reference to the agent in the description of the

https://doi.org/10.1017/S0265052509090062 Published online by Cambridge University Press

state of affairs that the agent has a duty to promote.[11] An example is the duty you have to give extra weight to the welfare of *your* friends when you are deciding how to allocate your own time, money, and other resources. A duty is agent-neutral if there is no essential reference to the agent in the description of the state of affairs that the agent has a duty to promote. An example is the duty to promote world peace.

As I said, complying with restrictions, respecting rights, and fulfilling special obligations are agent-relative duties. The restriction on physically harming others gives you a duty to avoid physically harming others yourself, not a duty to minimize the aggregate number of acts of physically harming. Thus, the agent-relative restriction on physically harming others forbids you to physically harm someone even if doing so would somehow prevent two other agents from physically harming two other victims to the same degree. The same is true of restrictions on stealing, damaging others' property, breaking promises, and telling lies.

It is also equally true of the duty not to infringe people's rights. The duty is to avoid infringing them yourself. The duty is not to minimize the aggregate number of infringements by anyone.

Likewise, special obligations are agent-relative. You are to give special weight to your friends and family when you are deciding how to allocate your time and other resources. This is not a duty to maximize aggregate fulfillment of everyone's special obligations.

There is one more agent-relativity that can conflict with maximizing aggregate social good. In order to see the rationale for this final kind of agent-relativity, we must first be clearer about what "social good" is supposed to mean. It obviously means the good of some multi-member set, the implicit contrast being between "social good" and "individual good." But does "social good" refer to the good of just a single society? There are difficulties about ascertaining the boundaries of societies, especially when societies overlap. These difficulties are especially troublesome if people have stronger obligations to members of their own societies than to outsiders. We might feel confident that we owe a special duty to members of our society, but not feel confident about where the boundary of "our society" is.

But consider now the aggregate social good calculated without regard to social, national, cultural, or other human boundaries. This aggregate social good is comprised of the good of every individual in the world counted impartially. This kind of aggregate social good can exert at least

[11] Here I am following David McNaughton and Piers Rawling, "Agent-Relativity and the Doing-Happening Distinction," *Philosophical Studies* 63, no. 2 (1991): 167–85. For influential earlier discussions, see Thomas Nagel, *The Possibility of Altruism* (Princeton, NJ: Princeton University Press, 1970), 90–96; Nagel, *The View from Nowhere* (New York: Oxford University Press, 1986), 153; and Parfit, *Reasons and Persons*, 143. Michael Ridge puts forward metaphysical objections to this account of agent-relative duties (Ridge, "Reasons for Action: Agent-Relative vs. Agent-Neutral," http://plato.stanford.edu/entries/reasons-agent/#4). I am not persuaded by these objections.

https://doi.org/10.1017/S0265052509090062 Published online by Cambridge University Press

https://doi.org/10.1017/S0265052509090062 Published online by Cambridge University Press

some moral pressure on everyone. That is, the good of someone geographically and culturally distant from us can give us a moral reason for doing something. For example, the suffering of any innocent person gives us some moral reason to stop it if we can.

At least a billion people in the world today live in severe poverty. Economists debate how much the richest sixth of the world's population can do for the poorest sixth, but let us accept that many things the richest sixth could do would hugely improve the prospects of the poorest sixth. Let us also observe that many sacrifices the rich could make for the poor would benefit the poor far more than they would cost the rich. Finally, let us observe that, under ordinary circumstances, most of the people in the richest sixth of the world's population will do little to help the world's poorest sixth.

Thus, a requirement on one of the people from the richest sixth to do what she can to maximize aggregate social good might be severely demanding. She would have to sacrifice an enormous amount before she reduced her welfare to the point where further sacrifices on her part would cost her more than they would benefit the very poor. In this way, an unlimited (e.g., act-utilitarian) requirement to maximize aggregate social good will be severely demanding in a world (like ours) where many people in a position to help do not do so, thereby increasing the burden on conscientious compliers with the requirement.

Excellent work has been done on this issue since Peter Singer's seminal essay "Famine, Affluence, and Morality."[12] Though not everyone agrees, many think that the demandingness of an unlimited duty to maximize aggregate social good seems to be a reason to limit this duty with some sort of agent-relative permission to refuse after some point to make further sacrifices for the sake of benefiting strangers. My own view is that the duty to do good for strangers should be qualified by a permission to refuse to make further sacrifices beyond some significant aggregate self-sacrifice.[13]

Actually, the appropriate permission is more complicated than just a permission to refuse to keep making sacrifices for the sake of others until

[12] Peter Singer, "Famine, Affluence, and Morality," *Philosophy and Public Affairs* 1, no. 1 (1972): 249–73. See also Samuel Scheffler, *The Rejection of Consequentialism* (Oxford: Clarendon Press, 1982); James Fishkin, *The Limits of Obligation* (New Haven, CT: Yale University Press, 1982); Shelly Kagan, *The Limits of Morality* (Oxford: Clarendon Press, 1989); Peter Unger, *Living High and Letting Die* (New York: Oxford University Press, 1996); Liam Murphy, *Moral Demands in Nonideal Theory* (New York: Oxford University Press, 2000); Tim Mulgan, *The Demands of Consequentialism* (Oxford: Clarendon Press, 2002); Garrett Cullity, *The Moral Demands of Affluence* (Oxford: Clarendon Press, 2004); and Tim Chappell, ed., *Moral Demandingness* (London: Palgrave, 2008).

[13] Here I have been most influenced by Garrett Cullity, "Moral Character and the Iteration Problem," *Utilitas* 7, no. 2 (1995): 279–89. See my *Ideal Code, Real World: A Rule-Consequentialist Theory of Morality* (Oxford: Oxford University Press, 2000), 147–74; my "Rule-Consequentialism and Internal Consistency: A Reply to Card," *Utilitas* 19, no. 4 (2007): 514–19; and my "The Demandingness Objection," in Chappell, ed., *Moral Demandingness.*

such sacrifices would cost the agent more than they would benefit others. Common-sense moral opinion holds that there is also a permission to make sacrifices for the sake of others where the benefit to them is less than the cost to oneself.[14] Common-sense morality does not insist upon impartiality between one's own interests and those of strangers, at least at the point of making day-to-day decisions. On the contrary, there is an agent-relative *permission* to give one's own interests either somewhat more, or the same, or somewhat less weight than the interests of others.[15]

Specifications of moral restrictions, rights, special obligations, and the agent-relative permission just mentioned of course provoke questions and concerns. There is a restriction on imposing physical harm on others (except in special circumstances), but what about emotional harm? How exactly are property rights to be defined? And how easily can they be overridden? What kinds or degrees of pressure can be put on someone to make an agreement, without undermining the moral force of that agreement? To whom are there special obligations? And what degree of favoritism is required or permitted? At what point has one made enough of a sacrifice to be permitted to refuse to go further? These are but some examples of the good questions posed about restrictions, rights, special obligations, and the agent-relative permission to give one's own interests either more or less weight than the interests of others.

Another good question is whether the restrictions, rights, special obligations, and permission are actually nonsense. Sometimes they are accused of being so. This is an accusation worth taking seriously.

The accusation that they are nonsense comes typically from the idea that the restrictions, rights, special obligations, and permission are not impartially defensible, but instead are biased toward the "haves" or at least toward the lucky. That line of attack used to get a sympathetic reception.[16]

More recently, many philosophers have entertained the idea that the credibility of the restrictions, rights, special obligations, and permission is higher than that of any impartial moral theory that rejects them. The idea takes the form of a rhetorical question: "By what right does abstract moral

[14] As far as I know, Michael Slote first called attention to this aspect of common-sense morality. See Slote, *Common-Sense Morality and Consequentialism* (London: Routledge and Kegan Paul, 1985), chap. 1; and Slote, *From Morality to Virtue* (New York: Oxford University Press, 1992), 13–16. More recently, Douglas Portmore has stressed this aspect in Portmore, "Position-Relative Consequentialism, Agent-Centered Options, and Supererogation," *Ethics* 113, no. 2 (2003): 303–32, at pp. 311–14; Portmore, "Consequentializing Moral Theories," *Pacific Philosophical Quarterly* 88, no. 1 (2007): 39–73, at p. 58; and Portmore, "Dual-Ranking Act-Consequentialism," *Philosophical Studies* 138 (2008): 409–27.

[15] I go into more detail about this in my "When Is Impartiality Morally Appropriate?" in John Cottingham, Brian Feltham, and Philip Stratton-Lake, eds., *Partiality and Impartiality: Morality, Special Relationships, and the Wider World* (Oxford: Clarendon Press, 2009).

[16] See, for example, Singer, "Famine, Affluence, and Morality"; and Kagan, *The Limits of Morality;* cf. Gilbert Harman, "Moral Relativism Defended," *Philosophical Review* 84, no. 1 (1975): 3–22.

https://doi.org/10.1017/S0265052509090062 Published online by Cambridge University Press

theory legislate to the moral sentiments?"[17] Given the strength of our convictions that there are restrictions on how one person should treat another, that people have rights that must not be infringed, that there are special obligations requiring certain kinds of partiality in one's private life, and that one is not required to make exorbitant sacrifices for the sake of strangers (expect perhaps in the most extreme circumstances), what abstract moral theory or conception of impartiality is so gripping that we should be willing to abandon all these convictions on its command? That question seems to me good enough to place a huge burden of proof on the shoulders of those who propose that the restrictions, rights, special obligations, and the agent-relative permission are all nonsense.

By no means am I rejecting moral impartiality as an ideal. Later in this essay (in Section IV), I will outline an argument for the view that moral impartiality can be reconciled with restrictions, rights, special obligations, and the permission. But the point I am accepting now is that our moral *epistemology* gives priority to the restrictions, rights, special obligations, and permissions over any proposed abstract principle of impartiality. In other words, we are more confident that there are moral restrictions, rights, special obligations, and permissions than we are in any proposed abstract principle of impartiality that rejects moral restrictions, rights, special obligations, and permissions. This is true even if we go on to embrace a moral structure that has an abstract impartial principle providing the moral foundation for all the more specific restrictions, rights, special obligations, and permissions. There are two different questions here: the question of what we are most confident about, and the question of what normatively underlies what. And these questions can get different answers.[18]

Let us turn from agent-relative reasons for refusing to maximize aggregate social good to agent-neutral considerations for refusing to maximize aggregate social good. One such agent-neutral consideration is equality of good. Suppose we take aggregate good to be a matter of aggregate welfare. And suppose we think the way to calculate aggregate welfare is to count everyone equally, in the sense that any benefit or harm to one person counts the same as the same size benefit or harm to anyone else. If aggregate social good is calculated this way, then one state of affairs could contain more aggregate social good than another despite the fact that in the first state of affairs some people have much more of a benefit

[17] Bernard Williams, *Moral Luck* (Cambridge: Cambridge University Press, 1981), x.

[18] See David O. Brink, "Common Sense and First Principles in Sidgwick's *Methods*," *Social Philosophy and Policy* 11, no. 1 (1994): 179–201. The theme is prominent in Robert Audi's moral philosophy. See, for example, Audi, "Intuitionism, Pluralism, and the Foundations of Ethics," in Walter Sinnott-Armstrong and Mark Timmons, eds., *Moral Knowledge?* (New York: Oxford University Press, 1996), 101–36; and Audi, *The Good in the Right: A Theory of Intuition and Intrinsic Value* (Princeton, NJ: Princeton University Press, 2004). It is a point I take up at the end of my "Intuitions and Moral Theorizing," in Philip Stratton-Lake, ed., *Ethical Intuitionism* (Oxford: Clarendon Press, 2002), 161–83.

https://doi.org/10.1017/S0265052509090062 Published online by Cambridge University Press

https://doi.org/10.1017/S0265052509090062 Published online by Cambridge University Press

than others do, and in the second state of affairs the benefit is distributed much more evenly. Concern for aggregate welfare can conflict with concern for equality of welfare.

Earlier I intimated that Pareto-superior outcomes are hard to resist. Yet, notoriously, concern for equality of welfare can lead one to oppose Pareto-improvements. Imagine a situation where someone could benefit from some change to the status quo and no one would lose from it. If the change would create or increase inequality of welfare, then a concern for equality of welfare would lead one to oppose the change, though it would make one person better off and no one worse off.

There are ways of incorporating concern for equality of distribution within the calculus of aggregate social good. The most widely discussed solution recently has been *weighted prioritarianism.* This theory attaches somewhat (but not infinitely) greater weight to a benefit to a worse-off individual than it does to the same size benefit to a better-off individual. Weighted prioritarianism, as I mean it, is *aggregative* in the sense that benefits to people are added together, with the possibility that a larger benefit to a worse-off person can be outweighed by smaller benefits to each of many better-off people.[19]

I am using the term "benefit" here, and throughout in this essay, not to refer to *instrumental means,* but rather to refer to what *constitutes* positive prudential value, that is, what *necessarily* adds to someone's welfare. Benefits for someone are what we care about insofar as we care about that person *for her own sake.* Likewise, I use the term "harm" to mean what *constitutes* negative prudential value.[20]

Weighted prioritarianism, then, is *not* the view that a unit of *material good or some other resource* has more moral importance when given to or taken away from someone who is worse off than it does when given to or taken away from someone who is better off. Weighted prioritarianism is not merely a way of coping with the diminishing marginal utility of resources and of other instrumental goods. Weighted prioritarianism holds that an addition to or reduction in the *welfare* of someone who is worse off (in terms of welfare) matters somewhat more than the same size addition to or reduction in the *welfare* of someone better off (in terms of welfare).

[19] An influential discussion of weighted prioritarianism can be found in Derek Parfit, "Equality and Priority," *Ratio* 10, no. 4 (1997): 202–21. For a ground-breaking discussion of the aggregative aspect of most versions of weighted prioritarianism, see Marc Fleurbaey, Bertil Tungodden, and Peter Vallentyne, "On the Possibility of Nonaggregative Priority for the Worst Off" (elsewhere in this volume).

[20] The concept of noninstrumental benefit and the concept of noninstrumental harm seem to me as basic as any in practical reasoning. For a different view, see Talbot Brewer, "Is Welfare an Independent Good?" (elsewhere in this volume). Stephen Darwall takes the concept of someone's welfare to be explained as what anyone who cares about her would rationally want for her *for her own sake.* See his *Welfare and Rational Care* (Princeton, NJ: Princeton University Press, 2002). In contrast, I take the concept of rationally caring about someone for her sake to be explained in terms of caring about her welfare.

Of course, an obvious question arises here: *How much* greater weight should be attached to benefits to the worse-off? I cannot see a principled way of arguing for any one clear answer to this question.

Nevertheless, there are some answers to the question that will enable weighted prioritarianism to circumvent some counterintuitive implications of utilitarianism. That is, weighted prioritarianism will be able to favor states of affairs that have somewhat less aggregate (impartially calculated) welfare but more equality, over other states of affairs with more welfare but less equality. This is because aggregate *weighted* good is greater when benefits go to the worse-off rather than to the better-off.

However, weighted prioritarianism has deep problems. I have already acknowledged the problem of determining how much greater weight should be attached to benefits to the worse-off. But there is a problem deeper than this, for there is a problem with prioritarianism no matter what level of extra weight it attaches to benefits to the worse-off.

Consider a view that I call *elitism.* Elitism takes different forms: aristocratic, talent-based, celebrity-based (as in the greater concern many people have for the welfare of this or that pop singer than for the welfare of the lady next door). For the sake of argument, focus on talent-based elitism, by which I mean a view that gives greater weight to benefits to the more talented than it does to the same size benefits to others.

Elitism is clearly morally wrong. Why is it wrong? The answer that immediately jumps to mind is that elitism is biased, or in other words not impartial. The way in which it is not impartial is obvious: it accords greater weight to additions to or reductions in the welfare of the more talented than it does to the same size additions to or reductions in the welfare of the less talented.

But if elitism is wrong because it is not impartial, then mustn't prioritarianism meet the same fate? Elitism and prioritarianism are structurally identical. Each of them gives extra weight to benefits for some people. The structure of elitism is what immediately disqualifies it. Since prioritarianism has the same structure, why isn't prioritarianism mistaken for the same reason?

One way of attacking this argument against weighted prioritarianism is to say that what is wrong with elitism is not that it is biased but rather that it is biased in the wrong way. Crudely, this attack claims that while a bias toward the talented is bad, a bias toward the worse-off is good.

I do not find this attack convincing. As I have indicated, what seems wrong to me about elitism is that it lacks a kind of fundamental impartiality that is desirable in a foundational moral principle. Admittedly, there are *kinds* of impartiality that elitism has. We can imagine the elitist saying, "Whoever the most talented are, benefits for those people count more than the same size benefits for others." Likewise, we can imagine the prioritarian saying, "Whoever the worse-off are, benefits for those people count more than the same size benefits for others." Nevertheless,

https://doi.org/10.1017/S0265052509090062 Published online by Cambridge University Press

there is another, deeper kind of impartiality that elitism manifestly lacks. This kind of impartiality consists in attaching equal importance to the welfare of each individual, that is, ascribing the same moral importance to an addition to, or reduction in, one person's welfare as is ascribed to the same size addition to, or reduction in, any other person's welfare. Prioritarianism lacks this kind of impartiality as much as elitism does.

My argument against weighted prioritarianism may be mistaken.[21] In a way, I hope it is mistaken, since, as I mentioned earlier, weighted prioritarianism seems to produce more intuitively attractive verdicts than utilitarianism produces. Still, a theory's claim to being an expression of fundamental impartiality is extremely important. I will say more about this in Section IV.

So far, I have discussed a number of different kinds of moral considerations that might lead us to oppose according aggregate welfare decisive weight. These were restrictions, rights, special obligations, the permission to give one's own interests greater or lesser weight than the interests of others, and concerns about equality (especially concerns about benefits to the worse off). There is one more consideration I want to discuss: desert.

We might sometimes face a choice between two alternative outcomes as follows. In one, aggregate welfare would be higher. In the other, although aggregate welfare would be lower, the distribution of benefits and harms would fit with people's deserts. Suppose we have some indivisible reward to distribute. Suppose Jack would benefit more than Jill from getting it, but Jill deserves it more. Or imagine that there is some inescapable cost to distribute, and the question is to whom to direct the cost. The cost might harm Jack more than Jill, but he might deserve the harm more than she does. In innumerable particular situations, maximizing aggregate welfare conflicts with giving people what they deserve.

Just as some of those sympathetic to egalitarian criticisms of utilitarianism have tried to answer these criticisms by making the calculation of aggregate good sensitive to equality, some of those sympathetic to the points above about desert have tried to answer these criticisms by making the calculation of aggregate good sensitive to desert. For example, Fred Feldman has advocated a "desert-adjusted" conception of the good, in which:

- Positive desert increases the intrinsic value of benefits.
- Negative desert decreases the intrinsic value of benefits.
- Positive desert increases the intrinsic disvalue of harms.
- Negative desert decreases the intrinsic disvalue of harms.

[21] Cf. Thomas Nagel, *Equality and Partiality* (New York: Oxford University Press, 1991), 65–70. I am grateful to Michael Otsuka for discussion about these matters.

https://doi.org/10.1017/S0265052509090062 Published online by Cambridge University Press

The increases and decreases operate, in Feldman's theory, to the point where some benefits have no intrinsic value, because they are undeserved, and some harms are not bad, because they are deserved.[22]

We may get an act-consequentialism yielding more plausible verdicts if this theory has a "desert-adjusted" conception of the good rather than a conception of the good making no reference to desert. But adjusting the good by importing into it values like desert weakens the explanatory power of the good. What gets built into the good cannot be justified by the good. I will illustrate this point again in the following section.

IV. The Structural Role of Aggregate Social Good

In this section, I contrast a range of theories about the relation of aggregate social good to the moral permissibility of acts.

One such theory is *act-utilitarianism,* the view that the moral permissibility of an act depends only on whether it maximizes expected aggregate welfare. Act-utilitarianism can advocate that agents have a "moral decision procedure" that focuses on restrictions, rights, special obligations, permissions, desert, etc. But act-utilitarianism holds that such things do not have any influence on the moral permissibility of acts.

What I call *Rossian moral pluralism* maintains that, on the one hand, aggregate social good is a fundamental moral consideration, but, on the other hand, there are other, equally fundamental moral considerations, such as nonmaleficence, trustworthiness, loyalty, desert, and self-improvement.[23] This view gives restrictions, rights, special obligations, permissions, and desert much higher status than act-utilitarianism does. Rossian moral pluralism holds that the moral permissibility of acts is determined partly by aggregate social good but also partly by restrictions, rights, special obligations, permissions, and desert. Indeed, Rossian moral pluralism typically holds that aggregate social good is a bit less important than other moral considerations, though it is at the same structural level as those others.

Another kind of theory that gives restrictions, rights, special obligations, permissions, and desert a role in determining the moral permissibility of acts is one kind of *agent-relative act-consequentialism.*[24] The kind of

[22] Fred Feldman, *Confrontations with the Reaper: A Philosophical Study of the Nature and Value of Death* (New York: Oxford University Press, 1992), 182–85; Feldman, *Utilitarianism, Hedonism, and Desert: Essays in Moral Philosophy* (New York: Cambridge University Press, 1997), 164–69.

[23] W. D. Ross, *The Right and the Good* (Oxford: Oxford University Press, 1930); Ross, *Foundations of Ethics* (Oxford: Oxford University Press, 1939). See also Thomas Nagel, "Fragmentation of Value," in Nagel, *Mortal Questions* (Cambridge: Cambridge University Press, 1979), 128–41.

[24] Many recent ideas about combining act-consequentialism with agent-relative restrictions and rights were inspired by Amartya Sen, "Rights and Agency," *Philosophy and Public Affairs* 11, no. 1 (1982): 3–39; Sen, "Evaluator Relativity and Consequential Evaluation,"

https://doi.org/10.1017/S0265052509090062 Published online by Cambridge University Press

https://doi.org/10.1017/S0265052509090062 Published online by Cambridge University Press

agent-relative act-consequentialism I have in mind holds that, while the permissibility of acts depends on the scenarios in which they occur, these scenarios are to be assessed differently from the points of view of different agents.

Suppose Jorge faces a strange dilemma: he can do one of only two things, namely, injure an innocent person himself, or not do so with the foreseen result that someone else will inflict similarly severe injuries on three other people. Common-sense morality holds that there is a restriction on injuring innocent others even when doing so would prevent a somewhat greater number of innocent people from suffering the same injury at the hands of some other agent. Agent-relative act-consequentialists can agree. They can say that, while the agent is to choose the act that promises to produce the best scenario, scenarios are to be assessed from the agent's point of view, and a scenario in which the agent injures one innocent person is worse from the agent's point of view than a scenario in which three innocent people are injured by someone else.

Agent-relative act-consequentialists may think that the really attractive feature of act-consequentialism is that it expresses the thought that doing what would produce the greatest expected good must be permissible. Agent-relative act-consequentialists might then try to build into their conception of the good everything needed to make their act-consequentialism come out with plausible verdicts about moral permissibility. I am prepared to accept that, with enough ingenuity, enough can be built into a conception of the good, especially if that good is agent-relative, for a version of act-consequentialism to come out with intuitively plausible verdicts about what is permissible. But I want to stress two objections to agent-relative act-consequentialism.[25]

One objection is that this kind of agent-relative act-consequentialism *cannot* take aggregate social good *to provide justification* for restrictions, rights, special obligations, permissions, and desert. In effect, this kind of agent-relative act-consequentialism must *postulate* restrictions, rights, special obligations, permissions, and desert, rather than providing a deeper and more abstract justification for them.[26] As I said at the end of the

Philosophy and Public Affairs 12, no. 2 (1983): 113–32; and Sen, "Positional Objectivity," *Philosophy and Public Affairs* 22, no. 2 (1993): 126–45. For more recent material, see Portmore's essays cited in note 14, though I stress that the views about goodness I attribute to agent-relative act-consequentialism are ones that Portmore's more recent essays have replaced with claims about what the agent has reason to prefer.

[25] As Frances Kamm has said to me in discussion, there may be better objections to agent-relative act-consequentialism than mine. I accept that observation, but at least these objections are mine.

[26] A common conception of the distinction between consequentialism and deontology is as follows. Consequentialism starts with a fairly minimal conception of the good (e.g., as welfare) and then derives the value of restrictions, rights, special obligations, permissions, and desert as instruments in the promotion of welfare; deontology gives restrictions, rights, special obligations, permissions, and desert a more foundational, less instrumental role. This common conception of the consequentialism/deontology distinc-

previous section, what gets built into the good cannot be justified by the good.

The other problem with agent-relative act-consequentialism is that it lacks an important kind of impartiality at the foundational level. Admittedly, agent-relative act-consequentialism can be impartially applied and can be framed in terms of universal principles. And, admittedly, partiality is often appropriate at the level of everyday decision-making. Nevertheless, there is enormous attraction in the idea that restrictions, rights, special obligations, permissions, and desert could be justified *from a wholly impartial point of view.* If we can find a moral theory that has no counterintuitive implications and that starts from complete impartiality, that theory will certainly be more attractive than any theory that either has counterintuitive implications (e.g., agent-neutral act-consequentialism) or, though it does not have counterintuitive implications, does not begin in impartiality (e.g., agent-relative act-consequentialism).

As I said, agent-relative act-consequentialists may think that the really attractive feature of act-consequentialism is that it expresses the thought that doing what would produce the greatest expected good must be permissible. In contrast, I think the really attractive feature of utilitarianism is its impartiality. Forced to choose between, on the one hand, giving up this impartiality and, on the other, rejecting the idea that doing what would produce the greatest expected good must be permissible, I have no hesitation in rejecting the idea that doing what would produce the greatest expected good must be permissible.

More generally, *rule-utilitarianism* has an impartial fundamental principle but rejects the idea that doing what would produce the greatest expected good must be permissible. This theory holds that impartial consequentialist assessment should be primarily of rules, habits of thought, and social practices. The theory then claims that which acts are morally permissible is determined by the rules, habits, and social practices selected because of their consequences. Rule-utilitarianism can leave the conception of the good fairly simple (impartial welfarism) and then emphasize the *indirect* connection between aggregate good and moral permissibility. If what I have argued elsewhere is correct, then rule-utilitarianism can provide justification for restrictions, rights, special obligations, permissions, and desert.[27]

The arguments of this section are summarized in table 1.

tion is threatened by any consequentialist theory that postulates, rather than derives, the value of restrictions, rights, special obligations, permissions, or desert. But my objection to agent-relative act-consequentialism is not that it plays havoc with the common conception of the consequentialism/deontology distinction. For an attack on the consequentialism/deontology distinction, see Adrian S. Piper, "A Distinction without a Difference," *Midwest Studies in Philosophy VII* (Minneapolis: University of Minnesota Press, 1982), 403–35.

[27] See chapters 6 and 8 of my *Ideal Code, Real World.*

https://doi.org/10.1017/S0265052509090062 Published online by Cambridge University Press

TABLE 1. *Social Good, Wrongness, and Deontological Elements*

	Gives aggregate social good the *sole* foundational role	Gives restrictions, rights, special obligations, permissions, and desert *some* role in determining wrongness	Takes aggregate social good, impartially calculated, to provide justification for restrictions, rights, special obligations, permissions, and desert
ACT-UTILITARIANISM	Yes	No	No
ROSSIAN MORAL PLURALISM	No	Yes	No
AGENT-RELATIVE ACT-CONSEQUENTIALISM	Yes	Yes	No
RULE-UTILITARIANISM	Yes	Yes	Yes

https://doi.org/10.1017/S0265052509090062 Published online by Cambridge University Press

V. Conclusion

This essay has explored what I think are some bad objections to aggregation, some good objections to it, and some bad ways of reacting to the good objections. Aggregation definitely has an important place in ethics, and this place might be a foundational one, together with impartiality. But aggregative considerations are neither omnipresent nor always overriding where they are present.

Philosophy, University of Reading

https://doi.org/10.1017/S0265052509090062 Published online by Cambridge University Press