Journal of Practical Ethics

Volume 3, Number 1. June 2015

CONTENTS

The Journal of Practical Ethics is available online, free of charge, at:
http://jpe.ox.ac.uk

Editorial Policy
The *Journal of Practical Ethics* is an invitation only, blind-peer-reviewed journal. It is entirely open access online, and print copies may be ordered at cost price via a print-on-demand service. Authors and reviewers are offered an honorarium for accepted articles. The journal aims to bring the best in academic moral and political philosophy, applied to practical matters, to a broader student or interested public audience. It seeks to promote informed, rational debate, and is not tied to any one particular viewpoint. The journal will present a range of views and conclusions within the analytic philosophy tradition. It is funded through the generous support of the *Uehiro Foundation in Ethics and Education*.

Copyright

ISSN: 2051-655X

Cost Effectiveness Analysis and Fairness [1]

F.M. KAMM

Harvard University

ABSTRACT

This article considers some different views of fairness and whether they conflict with the use of a version of Cost Effectiveness Analysis (CEA) that calls for maximizing health benefits per dollar spent. Among the concerns addressed are whether this version of CEA ignores the concerns of the worst off and inappropriately aggregates small benefits to many people. I critically examine the views of Daniel Hausman and Peter Singer who defend this version of CEA and Eric Nord among others who criticize it. I come to focus in particular on the use of CEA in allocating scarce resources to the disabled.

Cost Effectiveness Analysis (CEA) in medical care tries to maximize health benefits produced per dollar spent. Its use is recommended when society cannot afford every form of health care and must chose what to provide. Yet it is often taken as a truism that there can be deep conflicts between maximizing benefits and distributing fairly, in general. For example, the philosopher Robert Nozick imagined a "Utility Monster" [where utility is (roughly) experiential well being] such that for any resource up for distribution, one always produces more additional benefit at less cost if one gives the resource to the Monster rather than to others even though he is already

1. 1.This paper is a response to Daniel Hausman's "How Can We Ration Health Care Fairly and Humanely" as originally presented at "Bioethical Reflections: A Conference in Honor of Dan Brock," at Harvard Medical School Nov 22. 2014. All references to Hausman are to that paper. Hausman focused on Brock's discussions of the problems of fair chances, priority to the worse off, aggregation, and discrimination raised in several of his articles, including "Ethical Issues in the Use of Cost-Effectiveness Analysis for the Prioritization of Health Care." Hence, the order in which I discuss some issues in this paper follows the order in which Hausman chose to discuss Brock's work.

much better off than they. This would result in one person getting all the additional benefits while all others get none. This seems unfair.

CEA cannot result in this most extreme form of unfairness because of limits that result from how it calculates benefits. Each additional year of very healthy life is given a value of 1; no one can get more than 1 per year. Still, it is possible that only those who are already very healthy can achieve many additional years at a value of 1 at low cost if they are saved from an otherwise fatal bacteria. Maximizing health benefits per cost would imply helping them rather than people who are not as healthy and can achieve only fewer additional years at a value of less than 1 and at higher cost if they are saved from the same threat. This too seems unfair.

Why does it seem unfair to help only the Utility Monster and the healthy people? Fairness is about how one person is treated relative to another. This is by contrast with a notion such as justice which need not be comparative; that is, we could decide what justice requires that we give a person in virtue of his conduct independently of considering what anyone else is owed. Hence, we could treat someone justly in giving him what he is owed and thereby increase unfairness because we treat him justly while not treating anyone else justly. [For example, we might punish some who deserve this even if we cannot punish all who do. This case also shows that fairness is only one moral dimension on which we can evaluate how we treat people or states of affairs we produce; it is possible that we should sometimes override fairness to be just (or to achieve some other moral value).]

According to what measure shall we compare people to see if each is being treated fairly relative to others? Suppose we think that all that fairness requires in allocating benefits is that a certain amount of benefit be given the same value regardless of the person who will be benefited; no extra value should be assigned to a certain amount of benefit in person A rather than in person B. We could call this the Simple Standard. According to this standard, Nozick's Utility Monster and our CEA Bacteria case need not involve unfairness, for we could give the same value to a certain amount of benefit in the lives of everyone but it happens that more benefits can be produced in the Utility Monster and in the already healthy. These examples suggest that there may be more to fairness than the Simple Standard.

Indeed, there are different views about what fairness requires and deciding which is right is no easy matter. In this article, I will call attention to some different views of what fairness requires and consider whether, according to these views, fairness problems arise in the medical context as a result of using a version of CEA that

always emphasizes maximizing health benefits per dollar spent. Without pretending to settle the matter, I will raise some issues to consider. However, it is important to realize that even if there are problems with this version of CEA, this does not imply that it is never consistent with fairness to use some form of cost effectiveness evaluation. For example, it seems fair and right to treat one hundred people equally well with a cheap drug rather than an expensive one, other things equal. It would also be fair and right to use a drug with which we can save two hundred people rather than an equally costly one with which we can save only one hundred of these people.

1. Some think that when we cannot help everyone, fairness requires that people get a chance for medical care in proportion to their need for it, regardless of outcome in terms of CEA. For example, imagine that we can produce more cures per dollar if we treat six people with one fatal disease rather than only five people with another fatal disease. In this case the route to maximizing health benefits involves giving a life saving benefit to more rather than fewer equally needy people, by contrast with Nozick's Utility Monster. Hence, this case raises the question of whether it is unfair to save a greater number of people rather than give each group a chance to be saved in proportion to the need of each multiplied by the number of people in the group. Some think that fairness does not demand giving chances in proportion to need in the group but itself requires counting numbers of people, balancing one person of similar need (and, perhaps, expected outcome) against another and allowing the greater number to get the resource. On this view fairness does not require giving some chance to be helped to fewer people, when all suffer from equally serious problems.

Nevertheless, even on this view of fairness it may be right to give equal chances to be saved to two groups if they each contain the same number of people of equally sick people when we only have enough resources to treat one type of fatal disease. But proponents of CEA should see no reason to give equal chances if outcome per dollar would be the same. It is only if we take seriously the personal perspectives of each person, and so recognize that each person is not indifferent to whether he or someone else survives, that we see why fairness could sometimes require giving equal chances to different people even when their need and outcome are the same. (If we take seriously the perspectives of different people, we might even think it is wrong to deprive one person of his 50% chance to be treated merely because we would get a slightly better outcome if another person were treated.)

We have been focusing on cases in which each person needs a cure for a fatal

disease. However, if the original view of fairness that we considered (requiring chances in proportion to the strength of people's need) were true, someone who has a weak need for a scarce resource should get a small chance to get it. But does fairness really require giving a small chance to someone who needs the resource to cure his sore throat so that if against great odds he wins, someone else who needs the resource to save his life dies? Would we be overriding fairness in order to achieve a better outcome if we did not give the person with a sore throat a chance? I suspect not for in other cases achieving a better outcome would not lead us to override what fairness really requires. For example, suppose a doctor and a janitor both equally need a scarce life-saving medicine. If the doctor survives, he can save someone else's life from another illness, the janitor cannot. Although we could achieve a better outcome if we save the doctor (two people saved rather than one), this need not be sufficient reason to deny equal chances to the doctor and janitor in need of the scarce medicine. If we would not override fairness in this case to achieve the better outcome, this suggests that if we deny the person with the sore throat a proportional chance, we are doing so because we do not think fairness requires his having a chance, not because we are overriding fairness for the sake of a better outcome.

Hence, I have argued that while CEA does not necessarily contravene fairness in not giving chances to individuals in proportion to their need, it may fail to recognize an appropriate role for giving equal chances when need and outcome would be the same.

2. a. Another possible fairness concern about CEA is that it is indifferent to whether an equally cost effective benefit, such as relief from a certain amount of pain, goes to someone moderately ill or to someone severely ill. Some think fairness requires that the severely ill should be preferred. Indeed, it might be thought that fairness requires providing even a somewhat smaller benefit to those who are severely ill rather than a larger benefit to those who are moderately ill, holding cost constant. A "prioritarian" view of what fairness requires implies, roughly, that it is reasonable that the claim to benefits of those who are worse off should be weighed more heavily than those who are better off because it is right to give priority to improving the condition of a worse off person before improving someone who is already better off than

he is.[2] This view of fairness implies, contrary to the Simple Standard, that a given benefit in one person should sometimes have greater value than the same benefit in another person. (However, giving priority to the worse off is not the same as always caring for the worse off regardless of benefit that can be achieved.)

b. A third fairness concern about CEA is that providing minor, inexpensive health improvements (such as teeth fillings) in many people may be more cost effective than providing bigger, more expensive improvements (such as treating appendicitis) in a few, but fairness may require giving greater weight to the latter. This issue arises because CEA permits summing small benefits to each of many people to produce a large aggregate benefit that is then weighed against a smaller aggregate benefit composed of summing bigger benefits to a few people. The question is when it is fair to aggregate and weigh smaller benefits to some people against bigger benefits to others to decide how to allocate scarce resources.[3] This question about aggregation is sometimes related to the issue of not giving priority to the worse off, when the small benefits would go to many people already better off and the bigger benefits would to go to a few people more severely ill.

c. Let us consider these second and third concerns about the fairness of CEA in greater detail. With regard to the second, some think that willingness to help the severely ill even when this produces fewer benefits per dollar need not depend on a prioritarian conception of fairness. Rather it can reflect compassion for those in dire straits. Dan Hausman argues for this. On his view CEA is the reasonable, rational, and not unfair way to decide how to allocate medical resources but we sometimes override it because of compassion for the severely ill. On this view, compassion can conflict with reason (even if it does not always), and it is this compassion (rather than a reasonable view of fairness), that can lead us to help the severely ill when doing so conflicts with CEA.

One problem with this view is that it can seem fair and reasonable to give priority to treating those who are only moderately ill rather than those slightly ill even though compassion is not triggered for the former in the way it is for those in dire

2. There is a noncomparative view about giving priority to the worse off according to which the moral value of giving a benefit to someone varies with how well off in absolute terms that person is—the worse off, the greater the value. This view does not require comparing how well off someone is relative to others. I am focusing on the comparative prioritarian view in taking it to be an interpretation of fairness which is a comparitive value.

3. It is not unfair to aggregate and weigh small benefits to a few against the same small benefits to many others, other things equal. But I am concerned with aggregating and weighing smaller benefits against bigger ones.

straits. Similarly, we do not now feel greater compassion for twenty year olds who we know will die in thirty years at age 50 than for twenty year olds who we know will die at age 65. Yet we might still think it is morally right to invest in research that will buy five more good years for those who would die at 50 rather than in research that will buy 7 more good years for those who will die at 65, even if this conflicts with CEA. Presumably, this is because it seems reasonable to help people who would be worse off when they die, in having had shorter lives, than to help people who would be better off, in having had longer lives, even if they are not helped. This reflects a prioritarian conception of fairness.

Here is another problem with the view that it is compassion rather than a conception of fairness based on reason that sometimes conflicts with CEA. We often override compassion to do what is morally reasonable. For example, we may feel greater compassion for an incurably blind person who will also have to deal with a second problem if his arthritis is not treated than for a sighted person who will become (only) nearly blind if he is not treated. In this case, holding other factors constant, the blind arthritic will be the worst off person if he is not treated. Yet it seems morally acceptable and reasonable to cure the more severe condition of near blindness rather than the less disabling condition of arthritis, when we cannot do both. Suppose it is morally right to resist the call of compassion in this case. Then perhaps in other cases when we do *not* give up on helping the worst off person though helping him conflicts with CEA, it is because giving up would be contrary to reason and fairness rather than to compassion. This would imply that sometimes CEA is not the reasonable and fair approach.

Further support for the view that CEA does not necessarily coincide with what is reasonable and fair comes from considering the third concern mentioned above, whether it is always fair to aggregate small health benefits to many people and weigh the aggregate against the smaller aggregate of bigger benefits to fewer people, when the benefits to each group cost the same. Suppose that each of many people has a mild headache and is otherwise already much better off health wise than someone whom we can save from appendicitis. Suppose that none of the many people is a compassionate person and each would give up no more than the money for an aspirin that could cure his headache in order to help a dying person. But there are so many of these people with mild headaches that the aggregated harm of many headaches that would occur if each sacrificed the money for an aspirin is greater than the harm prevented in using the money to save the person with appendicitis. Though none of

the people is rescuing the one person because of compassion for him, presumably they would not change their individual aiding behavior on the grounds that the aggregated loss to all of them is so enormous by comparison to one person's loss of life. No one has to be compassionate in order to realize that it would be a bizarre mistake of reason to treat the very large aggregate of small losses to each of the many people as if it had the same moral significance as a very large loss to a single person that was suffered to prevent a smaller loss to another person.

Some suggest that the fair way to decide what to do in such cases, when the small harm (such as a headache) is occurring to each of many separate persons, is to compare, in a pair wise fashion, how much harm would be suffered and avoided by a severely ill person depending on whether he is helped with how much harm would be suffered and avoided by each of the the many depending on whether they are helped. Fairness is comparative but it requires comparing how we treat individual persons, one person at a time. Suppose that no one of the many will suffer anywhere near as great a loss as the single person would. Then if our view of fairness combines pairwise comparison with priority to the worse off, curing a headache in each of many people would never take precedence over curing even one much worse off person. A conception of fairness that involves these two components—pairwise comparison and prioritarianism—would justify the third concern with CEA.

By contrast, Peter Singer, a philosopher who supports CEA, believes it is morally correct to aggregate smaller individual benefits to better off people and weigh the aggregate against a bigger individual benefit to a worse off person. For example, in a *New York Times Magazine* article on rationing (Singer 2009), he considered how to compare the health benefit achieved in saving one person's life with curing a serious condition such as quadriplegia that does not threaten another person's life. He tells us to consider the trade-off each person would reasonably make in his own life between length of life and quality of life. Suppose every person (already disabled or not) would be indifferent between living ten years as a quadriplegic and living five years nondisabled. This seems to indicate that people take living as a quadriplegic to be half as good as living nondisabled.[4] Singer thinks that such data would show that using our resources to cure two quadriplegics is just as good as saving someone else's life when all three people would have the same life expectancy if helped (for example, ten

4. I say "seems" because it is possible that as the absolute number of years unparalyzed decreases (even if the ratio of unparalyzed to paralyzed years in the choice does not fall below ½), people would no longer be indifferent.

years). His reasoning seems to be that if someone would give up five out of ten years of his own life rather than be quadriplegic, that would justify curing one person's quadriplegia rather than saving someone else's life for five years, and the combined benefit of curing two people with quadriplegia would justify not saving the life of another person who would live for ten years.

Several things seem problematic about this reasoning. First, in the tradeoff between quality and quantity that a person might make in his own life, it is that person who benefits from the tradeoff. When we make tradeoffs between different people, the people who get the improved quality of life are not the same people who suffer the loss of more years of life. Trade offs between people may raise different moral issues than the trade off within one life. This is related to the point made earlier that fairness considerations arise when we take seriously that different people are not indifferent to whether benefits and losses fall in someone else's life or their own. Second, the conclusion that curing two quadriplegics who would live for 10 years anyway is equal to saving for ten years of life someone else who would otherwise die depends on aggregating the benefit to *two* people to weigh against the loss to the single person. We can see how problematic this is by considering the following example: Suppose that the trade-off test within one person's life showed that a small disability (e.g. a permanently damaged ankle) made life only 95% as good as a nondisabled life. This implies that a person would rather live 9 ½ years without the small disability than ten years with it. On Singer's view, this implies that we should cure one person's small disability rather than save someone else so that he can live for an additional ½ year. It also implies that we should cure the small disability in twenty-one people rather than save someone so that he can live for an additional ten years. This sort of problematic reasoning may have led to the rationing plan in Oregon many years ago in which resources were to be allocated to cap many people's teeth rather than save a few people's lives (for further discussion see Kamm 2007, chapter 2).

The concern with always aggregating small benefits can also be independent of concern to give priority to treating those more severely ill. For suppose all patients have the same disability. We have a choice of making very small improvements to the disability in each of a great many patients or providing a complete cure to one patient. It would not be morally unreasonable (and might not be unfair) to do the latter, for we then make a significant difference in this person's life rather than a barely perceptible

difference to each of many others. This is so even though aggregating the many barely perceptible differences creates an enormous total difference.[5]

3. Notice that in many of the cases we have considered, if a particular person is not helped, he will *certainly* suffer a great loss and be worse off than others who would certainly avoid only small losses and already be better off than he. But we may also consider the role of risk in deciding what is fair. We know that it can be reasonable for each individual to take a small risk of a great loss such as losing many years of life by dying if this is the price of achieving a high probability of lesser benefits. For example, someone might run a small risk of dying from an aspirin in order to get a high probability of relief from a nonlethal headache. If everyone in a community does this, in a large enough population it is certain that someone will die from an aspirin though each person had only a very small chance of dying. It seems morally permissible to allow individuals to expose themselves to such a small risk of the large loss of life for the sake of a small benefit. This is so even if we know that someone will die because he took the risk and because when he is dying there will be nothing we can do to save him. However, this need not mean that when it is still possible to save this person whose risk of dying has gone from small to certain, or when there are certain people who were always known to be at great risk of dying, fairness does not require aiding them if we can rather than protecting many people from each having a headache (One place in which I discuss this issue is Kamm 2008).

Suppose there are a few people who will certainly die unless we treat them. Should we treat them or rather use our scarce resources to stop a small risk of death to many others when it is certain that more than a few of these many will face certain death as well? This is the situation we may face when deciding whether to allocate scarce funds to combat AIDS by either treating those already ill or preventing future cases. Suppose fewer people overall will die from AIDS if the money is put into prevention than if it is put into treatment, and so prevention is most cost effective. But prevention deals with a small risk that each of many people has of dying while treatment deals with some people who will otherwise certainly die. (Notice that to make sure it is only the known probability of any given person dying whose relevance to an allocation decision we are judging, we should hold constant the time at which those already ill and those who will become ill would die. Otherwise, we may be judging the relevance to an allocation decision of sooner rather than later deaths, not greater or smaller probabilities each person has of dying.) It is true that without

5. Larry Temkin has emphasized this point.

prevention more people who once had only a small risk of getting AIDS will eventually face certain death. But at the time we must decide how to spend funds there are certain people who already have a known prognosis of certain death if they are not treated now whereas it is not the case that there are already some people in the larger group who have a known prognosis of certain death if prevention measures are not taken now. Rather there are many people each of whom has a small chance of being a person who will face certain death. So we must choose between helping those who will certainly die and those who each have a small chance of facing death.[6]

On at least one view, fairness seems to require helping those with a higher individual risk of dying as ascertained by a pair wise comparison of the risks each person faces at the time we must allocate resources. This is because if we engage in pairwise comparison, we could justify to each person with the low risk of death not helping him and helping the person who will certainly die instead. By contrast, it seems we could not justify to the person who will certainly die leaving him to help each of those who have a small chance of death. On this view, it is individuals' comparative risk at the time we must allocate, not the ultimate outcome in which more rather than fewer people will die, which should lead us to allocate the money to the less cost effective treatment policy. This is so even if fairness requires saving the greater number of people when these are all people who face certain death if not helped.

4. Another possible fairness concern with CEA is that it might involve discrimination against those who are poor or disabled as it may cost more to treat these people by contrast to the rich or nondisabled, and the health benefits achieved may also be less. Peter Singer relies on CEA when he argues that if we accept that disability can make a person's life less good health wise, and we want to maximize the health benefits we get with our resources, we should save the life of a nondisabled person rather than someone whose disability cannot be cured, other things equal. The only alternative to this, he says, is to deny that disability per se makes someone's life not as good health-wise, and then there would be no reason to allocate resources to cure or prevent disabilities which seems wrong. However, there are other alternatives, I think.

I agree that understanding the issue of disability and allocation of scarce resource should not depend on accepting the view that disabilities make little difference to the quality of life. For if we hold this view, we may see little reason to invest

6. This analysis of the case is argued for by Norman Daniels (2012) and Johann Frick (2013 and unpublished) though they may not hold time of death constant.

in curing disabilities. We should also recognize that a satisfied "mood" is not the sole measure of the goodness of one's life, independent of objective capacities. Consistent with all this one proposed response to a view like Singer's is offered by Eric Nord, Norman Daniels and Mark Kamlet in their article (2009).[7] They think it is important to distinguish two different questions. The first is: "Is a health state one we would prefer to cure?" Suppose we answer yes. A second question arises if we have this health state and it cannot be cured but life is still worth living, and we also have a life-threatening treatable condition but the medicine is scarce. The second question as stated by the authors is: "Should we defer to those who can be restored to more complete health than we can because they lack the untreatable condition?" The authors say we can reasonably answer 'yes' to the first question—we would prefer a cure to the health state—and 'no' to the second question. They do not say what explains the reasonableness of these responses. Given the way their second question is phrased, it might be thought that one simple explanation is that the person with one untreatable condition does not have a duty to defer because he does not have a moral duty to sacrifice what is very important to him (his life) to produce the outcome that would be considered best from an impartial point of view. The fact that this view, which is standardly held by those who reject consequentialism, might explain the consistency of the first and second answers suggests to me that the second question as phrased by Nord, Daniels and Kamlet is the wrong one to pose if we want to get to the heart of the issue in allocating scarce resources.

This is because it should be an impartial distributor who is allocating the resources, not a candidate for the resource and the mere fact that a candidate need not defer to another candidate does not mean that the impartial distributor must give these people the same chance of treatment. Analogously, someone need not give up his medicine that will save his leg so that someone else may use it to save his own two legs. But if the drug is publicly owned and to be distributed by an impartial agent, he should prefer to help the person who would otherwise lose two legs. So a crucial issue in dealing with Singer's CEA-inspired view is whether, if an impartial distributor says yes to the first question, this distributor should also decide to treat the person who can be restored to more complete health. In my own past work (first in Kamm 2004), I have been interested in the answer to this question (which I shall refer to as the impartial question).

7. Nornan Daniels brought my attention to what was said in this article in his commentary at a panel in February 2013.

Suppose I am the impartial distributor. When I imagine a case in which someone has a paralyzed finger, I can see that this can make life not as good in a small way, other things equal, and give us a reason to fix the disability. Hence, my answer to the first question is yes. But when I consider whether to save someone's life from pneumonia when I can only save one person, the fact that one of the people I would save has a missing finger and the other has all his fingers should, I think, make no difference to whom I choose, given the important benefit that is at stake for each person and that each person desires to be the one to live. Hence, I should answer "no" to the impartial question. One explanation I have suggested for this is that a factor (such as a missing finger) could give us a reason to act in one context (curing it) while it is an irrelevant consideration in another context where the action in question (saving a life) is different. So it was an instance of what I call "contextual interaction."[8]

However, suppose a nondisabled candidate for a scarce life saving surgery will live for twenty years after it and a disabled candidate will live for only two years because his disability interferes with doing exercise after surgery. A large difference in length of life in the outcome might be a morally relevant difference between the candidates, and that might make it not be unfair to do the surgery on the nondisabled candidate, given that each will die if not helped. This is so even if we would be using an effect of the undeserved or unjustified disadvantage in the disabled as grounds for choosing to impose further disadvantage on him.

Furthermore, consider another case about which we ask the first question and the impartial question. Suppose we agree paraplegia is a deficit that we should prefer to cure. Now imagine two people with paraplegia who each need to be saved from fatal heart disease. The only difference between them is that in one of the people the scarce heart disease medicine will also cure his paraplegia. This is also a case in which one candidate has an untreatable condition (paraplegia) and a treatable one (heart disease) and another candidate lacks the untreatable condition (because his paraplegia can be cured). I suggest that it might be right for the impartial distributor of a scarce resource to choose to save the candidate whose paraplegia will also be cured rather than the other candidate. However, I think that this is not simply

8. One counterargument to this that Singer also gives is meant to show that it is reasonable to connect the answer to the first question to an answer to the impartial question. This is because, he thinks, the morally right way for an impartial allocator to decide is determined by what any person would decide reasoning about his possible future state when he is ignorant of which person he will be. I do not think this is correct and argue against it in Kamm (2013). But it is useful to see an argument, aside from maximizing good outcomes, that has been thought to connect a yes answer to the first question to a yes answer to the impartial question.

because being alive with paraplegia is worse than being alive with full mobility (other things equal) and CEA rates a treatment as more effective if a person is saved to a life of higher rather than lower quality. Rather, it has something to do with how bad paraplegia is and with our medical procedure causing the paraplegia cure. For my sense is that it's not being unfair to treat the second person whose paraplegia we can cure does not imply that in a different case it would be fair to treat the heart condition of someone who is not paralyzed independently of what we do rather than treat a permanently paralyzed person who also has the treatable heart condition.[9]

These two heart cure cases suggest that a possible problem with CEA is that it does not distinguish the first case, in which our treatment is more cost effective in one candidate because it saves a life and also *causes* the change in disability status, from the second case in which our treatment is more cost effective in one candidate only because it saves someone who is already nondisabled independent of our doing anything to cure him of disability. However, the two heart cases do not suggest that if each of two people has a paralyzed finger and a scarce life-saving drug that each needs to save his life will unparalyze the finger in only one of the people, that we should give that person the medicine. That a condition is one that we would prefer to cure does not mean that it is serious enough in itself that our being able to cure it in one person but not another should make a difference to whom we give a drug needed by each for a much more serious condition.

5. Conclusion: In this article, I have considered several views about what fairness requires and allows, in conditions of certainty and risk, and how CEA understood in its strongest form may conflict with fairness. It has not been my aim to decide which conception of fairness is correct or to decide how important fairness is relative to other moral considerations. Nor has it been my aim to deny that cost effectiveness should sometimes play a role in allocating scarce resources. However, if the value of maximizing good outcomes relative to cost is neither a preeminent value nor necessarily fair, there are bound to be moral questions about limits on the use of CEA that will need to be resolved.

Acknowledgements: I am grateful to Julian Savulescu and a reader for the Journal of Practical Ethics for comments.

9. I consider this issue in more detail in (Kamm 2013), a revised version of (Kamm 2009)

REFERENCES

Daniels, Norman. "Reasonable Disagreement about Identified vs.Statistical Victims " *Hastings Center Report*, 42, no.1 (2012), 35- 45

Frick, Johann. "Uncertainty and Justifiability to Each Person: Response to Fleurbaey and Voorhoeve " in *Inequalities in Health* ed. N. Eyal et al (New York: Oxford University Press, 2013)

———. "Treatment and Prevention of HIV/AIDS and the Question of Identified vs.Statistical Lives" (unpublished)

Kamm, Frances M. "Deciding Whom to Help, Health Adjusted Life-Years, and Disabilities," in *Public Health, Ethics and Equality*, eds. S. Anand, F. Peter, and A. Sen (NewYork: Oxford University Press, 2004)

———. *Intricate Ethics* (New York: Oxford University Press, 2007)

———. "Should You Save This Child? Gibbard on Intuitions, Contractualism, and Strains of Commitment." in *Reconciling Our Aims* , ed. A. Gibbard (New York: Oxford University Press, 2008)

———."Aggregation, Allocating Scarce Resources, and the Disabled," *Social Philosophy and Policy* 26 (Winter 2009), 148-97

———. "Aggregation, Allocating Scarce Resources, and the Disabled," *Bioethical Prescriptions: To Create, End, Choose, and Improve Lives* (New York: Oxford University Press, 2013)

Nord E, Daniels N, and Kamlet M. "QALYs: some challenges." *Value in Health*. (Mar 2009) 12 Suppl 1;S10-5

Singer, Peter. "Why We Must Ration Care," *New York Times*, July 19, 2009

The Elements of Well-Being

BRAD HOOKER

University of Reading

ABSTRACT

This essay contends that the constitutive elements of well-being are plural, partly objective, and separable. The essay argues that these elements are pleasure, friendship, significant achievement, important knowledge, and autonomy, but not either the appreciation of beauty or the living of a morally good life. The essay goes on to attack the view that elements of well-being *must* be combined in order for well-being to be enhanced. The final section argues against the view that, because anything important to say about well-being could be reduced to assertions about these separable elements, the concept of well-being or personal good is ultimately *un*important.

CLARIFICATIONS

Like most philosophers, I will take the term 'well-being' to be synonymous with 'welfare', 'personal good', and 'individual utility'. Contributions to well-being I refer to as benefits or gains. Subtractions from well-being I refer to as harms, losses, or costs. The elements of well-being are whatever constitutes benefits, that is, contributions to well-being.

Absolutely essential is the distinction between non-instrumental value, which is sometimes called final value, and instrumental value. Examples of things with merely instrumental value are money, medicine, and sleep. This paper focuses on non-instrumental value. When I refer to contributions to well-being, I mean *non-instrumental* contributions, that is, things that are good for us in their own right as opposed to good only because they are means to other things. The main focus of the paper is on

the question of what constitutes non-instrumental contributions to a person's well-being. Definitely, all of the values I will be discussing do have instrumental value. But I will be focusing on these values not for their instrumental benefits but as putative elements of well-being.

HEDONISM

Hedonism is the theory that well-being consists in pleasures minus pains. Pleasures are experiences found attractive solely because of their experiential quality, rather than for other reasons. Pains are experiences found aversive solely because of their experiential quality, rather than for other reasons.[1]

The focus on experiential quality brings out a defining feature of pleasures and pains, namely that they are *introspectively discernible* (which is not to say that they are *actually* discerned). Imagine someone who fails in the central project of her life but never finds out about this failure. An example might be the amateur sleuth who spent the last five years of her life trying to discover how and why the child Madeleine McCann disappeared. The sleuth died thinking that she had made the crucial discovery that solved the case. But in fact her 'discovery' turned out to be quite mistaken. Because she didn't find out that she failed, her pleasures were what they would have been had her project instead been a success. Hedonists hold that the failure of a life project does not, in itself, reduce the person's welfare. Hedonists think that a person's welfare is determined solely by how this person's life feels from the inside. How her life feels from the inside may depend in part on whether she *believes* her desires have been fulfilled. How her life feels from the inside does not necessarily depend on whether her desires *really* have been fulfilled. (For recent defences of hedonism, see Feldman, 2004, and Crisp, 2006, ch. 4.)

I have contended that introspective discernibility is essential to pleasure and that success in one's projects is not. This is true whether the project is relatively discrete, such as finding out how and why a small child suddenly disappeared, or much more complex and general, such as the goals of having lots of good friends and of being knowledgeable about science, history, and metaphysics and of creating things of enduring value. Consider someone who believes that he has enough good friends and that he is knowledgeable about science, history, and metaphysics and that he has

1. Here I am especially grateful to anonymous reviewers for wording.

created things of enduring value. This person is likely to feel some satisfaction with his life.

Perhaps this sort of satisfaction is the most important kind of pleasure (Sumner, 1996, ch. 6). Nevertheless, getting this kind of pleasure is possible even if one is deluded about whether one's desires for good friends, for knowledge of science, history, and metaphysics, and for creative success have actually been fulfilled. Feeling satisfied with one's life is compatible with delusion about pretty much everything except whether one feels satisfied with one's life.

DESIRE-FULFILMENT

Another main view of welfare holds that a person's well-being is constituted by the fulfilment of his or her desires, whether or not the person knows the desires have been fulfilled. This view is often called the desire-fulfilment (or preference-satisfaction) theory of well-being.

The main argument in favour of the desire-fulfilment theory over hedonism is that many people's self-interested concern extends beyond their own pleasures and pains, enjoyments and frustrations (Nozick, 1974, esp. p. 43). For example, many people have stronger self-interested concern for knowing the truth (especially about whether their other desires are fulfilled) than for blissful ignorance.

The main argument against the desire-fulfilment theory is that some desires are so wacky that their fulfilment would not itself constitute a benefit for the people who have them (even if whatever associated pleasure these people derived from believing their desires were fulfilled *would* constitute a benefit for them). Imagine someone who wants a saucer of mud, or to count all the blades of grass in the lawns along a street, or to turn on as many radios as possible (Anscombe, 1958, p. 70; Rawls, 1971, p. 432; Quinn 1993, p. 236). Suppose this person wants these things *for their own sakes,* i.e., non-instrumentally. Fulfilment of such desires in itself would not be of any benefit to this person, we intuitively think.

OBJECTIVE LIST THEORY

A third theory of welfare agrees with hedonism that pleasure constitutes a benefit. Where this third theory departs from hedonism is over the question of whether there is only one element of well-being or more than one. The third theory

claims that other things can also constitute benefits—for example, knowledge of important matters, friendship, significant achievement, and autonomy. Derek Parfit (1984, pp. 493–502) dubbed this theory the 'objective list theory', but often the name is shortened to the 'list theory'. According to this objective list theory, a life contains more welfare to the extent that it contains pleasure, knowledge of important matters, friendship, significant achievement, and autonomy. A life full of pleasure and fulfilment of desires for things other than the goods just listed could still be of low quality precisely because it lacked the goods just listed. (For discussion, see Griffin, 1986, pp. 29–35, 58–72; Crisp, 1997, ch. 3.)

What makes one achievement more significant than another? Thomas Hurka (1993, chs. 8–10; 2011, ch. 5) argues persuasively that extended and difficult achievements are more significant than narrow and easy ones. Admittedly, a narrower and less difficult achievement might benefit you more than one that is more extended and difficult, because the narrower one brings you greater pleasure or because it helps develop your friendships or because you learn more from it. In other words, when instrumental as well as intrinsic value is considered, a narrower and less difficult achievement can be on balance more beneficial to you than a wider and more difficult achievement. But when we ignore the instrumental benefits of different achievements, we should conclude that extended and difficult achievements are more significant than narrow and easy ones.

We might likewise follow Hurka (1993, chs. 8–10; 2011, ch. 4) in holding that extended and explanatory knowledge is better than narrow and shallow knowledge. For example, knowledge of the basic truths of physics or biology or metaphysics is more important than knowledge of the batting averages achieved by the middle-ranking players on a particular team in a particular month thirty-three years ago. But knowledge about yourself or things closely connected to you can sometimes constitute a larger benefit to you than would more general knowledge about things with no special connection to you. Knowing important facts about yourself—having self-knowledge—is a more important element of your well-being than knowing general truths about physics or biology or metaphysics or other people. For example, knowing your own failings is more important than knowing other people's failings. On the other hand, knowing that something is true not only of you but also of everyone else would constitute a larger benefit than knowing merely the truth about yourself.

Even more contestable than which kind of achievement or knowledge is most valuable as an element of well-being is the question of exactly what comprises auton-

omy. Does autonomy consist merely in having one's actions be guided by desires that one desires to have? Or does autonomy consist in having one's decisions be guided by one's own value judgements? Or does autonomy require that one's value judgements be themselves autonomously produced? Or does autonomy require that one's value judgements be at least minimally sensible?

These questions are fascinating but, alas, too difficult to address here. Hence, I must simply assume an answer. This is that someone's life contains autonomy to the extent to which she has a variety of important options to choose among, her choices reflect her value judgements, and her value judgements are at least minimally reflective (i.e., she has at least once considered them rather than merely always accepted them without consideration). If Jack severely constricted Jill's set of important options, or if he controlled her choices by controlling her value judgements, she would lack autonomy. The same would be true if a brain injury or mental illness controlled her value judgements or prevented her from being able to assess them.

Now, how can we ascertain whether any given putative good is an item on the objective list? We must run the following kind of thought experiment. We imagine two possible lives for someone that are as much alike as possible except that one of these lives contains more of some candidate good than the other. We then think about whether the life containing more of the candidate good would be more beneficial to the person living it than the other life. If the correct answer is no, then definitely the candidate good in question is not an element of well-being. On the other hand, if the correct answer is instead that the life with more of the candidate good is more beneficial, then we inquire what is the right explanation of this life's being more beneficial. One possible explanation is that the candidate good in question really is an element of well-being.

PLEASURE

Here is an illustrative example. We imagine two possible lives for someone that are as much alike as possible except that one of these lives contains a larger amount of innocent pleasure than the other. We are trying to hold everything equal as much as possible with the single variable being the amount of innocent pleasure in the two possible lives. We then think about whether the life containing the larger amount of innocent pleasure would be more beneficial to the person living it than would be the as similar as possible life with a smaller amount of innocent pleasure. If the correct

answer is that the life containing the larger amount of innocent pleasure would *not* be more beneficial to the person living it than would be the as similar as possible life with a smaller amount of innocent pleasure, then innocent pleasure is not an element of well-being. On the other hand, if the correct answer is that the life with a larger amount of innocent pleasure *is* more beneficial, then we need to inquire what is the right explanation of this life's being more beneficial. The explanation that suggests itself is that innocent pleasure is indeed an element of well-being.

For that possible explanation to be correct, rival possible explanations must be mistaken. Perhaps the leading rival possible explanation is that, although by hypothesis the two lives being compared are as much alike as possible with the exception that one includes a larger amount of innocent pleasure than the other, the fact that one of these possible lives contains a larger amount of innocent pleasure brings with it differences in the levels of *other* goods and these differences are what account for the superiority of one possible life to the other. In short, although our thought experiment was supposed to isolate one variable, the rival possible explanation claims that other variables are not only ineliminable but also pivotal.

Here is an example of such a rival explanation. This explanation begins with the proposal that the life with the larger amount of innocent pleasure must also have contained a larger amount of significant achievement or friendship or important knowledge or autonomy, as sources of the extra innocent pleasure. This rival explanation then adds that what makes the life with the larger amount of innocent pleasure more beneficial to the person who lives it than the life with a smaller amount of innocent pleasure is not the extra innocent pleasure but instead the larger amount of significant achievement or friendship or important knowledge or autonomy.

This rival explanation starts from a false supposition—namely, that the life with the larger amount of innocent pleasure *must* also have contained a larger amount of significant achievement or friendship or important knowledge or autonomy, as sources of the extra innocent pleasure. This supposition is false because the extra pleasure might have come from insignificant achievement or unimportant knowledge or false beliefs or the satisfaction of physiological urges. The source of innocent pleasure can be trivial or misconceived or merely physiological. There is no necessity that source of innocent pleasure is itself something valuable, much less an element of well-being.

We can conclude, then, that innocent pleasure is definitely an element of well-being. This is the best explanation of why a life containing a larger amount of in-

nocent pleasure would be more beneficial to the person living it than another possible life as much as possible like the first one except that this life contains a smaller amount of innocent pleasure.

I am aware that, to many people, the thesis that innocent pleasure is an element of well-being seems completely obvious. However, there are some people who do not find innocent pleasure obviously valuable. My argument above is aimed at these people. Later in this essay, there is an argument aimed at people who think innocent pleasure is valuable only if obtained in the course of a worthwhile activity.

SIGNIFICANT ACHIEVEMENT

The structure of the argument above can be applied to other candidate elements of well-being. Let us apply it to significant achievement. We imagine two possible lives for someone that are as much alike as possible except that one of these lives contains a larger amount of significant achievement than the other. We are trying to hold everything equal as much as possible with the single variable being the amount of significant achievement in the two possible lives. We then think about whether the life containing the larger amount of significant achievement would be more beneficial to the person living it than the life that contains a smaller amount of significant achievement but is otherwise as similar as possible. If the correct answer is no, then significant achievement is not an element of well-being. However, the correct answer seems to me to be that the life with a larger amount of significant achievement is more beneficial to the person who leads that life than the life that contains a smaller amount of significant achievement but is otherwise as similar as possible.

Now, what is the best explanation of this life's being more beneficial? One possible explanation is that significant achievement is indeed an element of well-being. The rival possible explanation starts from the supposition that the life with the larger amount of significant achievement must also have contained a larger amount of innocent pleasure or friendship or important knowledge or autonomy. From this supposition, the rival explanation infers that what makes the life with the larger amount of significant achievement more beneficial to the person who lives it is not the extra significant achievement but instead the larger amount of pleasure or friendship or important knowledge or autonomy.

Lives containing a larger amount of significant achievements often do also have more pleasure, friendship, important knowledge, or autonomy in them than they

would have had if they had contained a smaller amount of significant achievement. But this certainly is not always true. Sometimes people sacrifice pleasure, friendship, and important knowledge for the sake of pursing significant achievement. Indeed, obsession about a goal can be instrumental to achieving it but at the same time in conflict with obtaining pleasure, alienating to friends and potential friends, and a blinker to information not relevant to the goal. For such reasons, there is no necessity that a larger amount of significant achievement correlates perfectly with a larger amount of pleasure, friendship, or important knowledge. This is especially obvious in cases where the person who made a significant achievement never found out about it, and thus could not have gained lots of pleasure from knowing about the achievement.

So what is the correct explanation of the fact that the life containing the larger amount of significant achievement would be more beneficial to the person living it than the life that is as similar as possible except that it contains a smaller amount of significant achievement? The correct explanation cannot be that the life with the larger amount of significant achievement *must* also have contained a larger amount of innocent pleasure or friendship or important knowledge or autonomy. The correct explanation is instead that significant achievement is an element of well-being.

IMPORTANT KNOWLEDGE

The same kind of argument can be run for concluding that important knowledge is an element of well-being. Imagine that two lives are as much alike as possible except that one of those lives has important knowledge and the other life does not or one life contains a considerably larger amount of important knowledge than the other. The life containing no or a considerably smaller amount of important knowledge is worse for the agent than a life as much as possible like that one except that it contains at least some or a considerably larger amount of important knowledge.

Again, we have to ask, what is the best explanation of the greater benefit in the life with a larger amount of important knowledge? One possible explanation is that important knowledge is indeed an element of well-being. The rival possible explanation starts from the supposition that the life with the larger amount of important knowledge must also have contained a larger amount of innocent pleasure or more friendship or a larger amount of significant achievement or greater autonomy. From this supposition, the rival explanation infers that what makes the life with the larger amount of important knowledge more beneficial to the person who lives it is not the

extra important knowledge but instead the greater pleasure or friendship or significant achievement or autonomy.

We should not accept the supposition that the life with the larger amount of important knowledge must also have contained a larger amount of innocent pleasure or friendship or significant achievement or autonomy. Sometimes important knowledge reduces rather than increases innocent pleasure. Sometimes important knowledge harms friendships. Sometimes important knowledge is not a significant achievement because it was not something that was pursued and thus not an achievement at all. Hence the life with the larger amount of important knowledge *might not* also contain greater innocent pleasure or friendship or significant achievement or autonomy.

Thus, the best explanation of the fact that a life containing no or a considerably smaller amount of important knowledge is worse for the agent than a life as much as possible like that one except that it contains a larger amount of important knowledge *cannot* be that the life with the larger amount of important knowledge must also contain greater innocent pleasure or friendship or significant achievement or autonomy. The best explanation must instead be that important knowledge is an element of well-being.

AUTONOMY

Concerning autonomy, we can try an argument with the same structure as the arguments above. We imagine two possible lives for someone as similar as possible except that one contains more autonomy and the other less. Then we ask which of these two possible lives is more beneficial to the person who lives it. The autonomous life seems better. Since we have imagined that the two lives are as equal as possible in terms of the other elements of well-being, we minimize the extent to which the superiority of the more autonomous life can be explained by the instrumental value of autonomy.

Again, we face the objection that there are multiple, though correlative, variables here. The objection is that, if one possible life contains greater autonomy than another possible life that is otherwise as similar as possible, the possible life with greater autonomy in it *must* also contain greater pleasure or knowledge or friendship or achievement than the other life. From this supposition, the objection infers that the greater pleasure or knowledge or friendship or achievement, and not the greater autonomy, is what makes this life better.

However, it just is not true that if one possible life contains greater autonomy than another life that is otherwise as similar as possible, then the possible life with greater autonomy in it must also contain greater pleasure or knowledge or friendship or achievement. Greater autonomy might lead to some successes but also, of course, to some failures—some missed opportunities for pleasure, some lack of knowledge, some ruined friendships, and some unsuccessful projects. So, on balance, greater autonomy might not lead to greater pleasure or knowledge or friendship or achievement. So there is nothing to prevent our imagining two possible lives that are equal in terms of pleasure, friendship, achievement, and knowledge and yet one of these lives contains more autonomy than the other.

We thus ask which of these two possible lives is more beneficial to the person who lives it. The more autonomous life seems better. Since we have imagined that the two lives are as equal as possible in terms of the other elements of well-being, the best explanation of the superiority of the more autonomous life in this comparison is that autonomy is an element of well-being.

APPRECIATING BEAUTY

Experiencing something as beautiful can definitely be instrumental to other benefits. Most obviously, experiencing something as beautiful can produce pleasure, even ecstasy. Experiencing something as beautiful can also lead to other goods, such as love and knowledge. But is experiencing something as beautiful a non-instrumental good such that a life containing such appreciation must be pro tanto better than a life without?

Well, are judgements of beauty like judgements of taste, i.e., merely subjective? Whether a food is delicious or not is merely subjective. For example, if you judge pears to be delicious and I do not, neither of us need be mistaken. If beauty is like deliciousness, then while you can be correct about whether you find something's aesthetic qualities attractive, you cannot be correct about whether these qualities *really* are attractive or about whether they *ought* to attract. If beauty is like deliciousness, then what would make a possible life in which you find more beauty better for you than a possible life in which you find less beauty would be the additional pleasure or friendship that the extra beauty would bring you. If beauty is like deliciousness, then appreciating beauty is not itself an element of well-being.

If beauty is not like deliciousness but is instead an objective value, then you can

be correct not merely about whether certain qualities attract you and others but also about whether certain qualities *really* are attractive or about whether they *ought* to attract. In that case, you can have knowledge of aesthetic properties. If this is correct, perhaps we should classify appreciation of beauty as a kind of important knowledge. If appreciation of beauty is a kind of important knowledge, then we have grounds for holding that appreciation of beauty is an element of well-being—under the heading of important knowledge.

LIVING A MORALLY GOOD LIFE

Let us now turn to the question of whether living a morally good life is an element of well-being. Even if it is not, living a morally good life is of course *morally* good. And living a morally good life might be what there is strongest reason to do even where living a morally good life involves self-sacrifice. We can be interested in the question of whether living a morally good life is an element of well-being even if we are committed to sacrificing our own good either for the sake of benefiting others or because moral restrictions get in the way of doing what is most beneficial to oneself.

Living a morally good life is rewarding in terms of the other elements of well-being. Living a morally good life definitely constitutes a significant achievement.[2] And living a morally good life can bring pleasure and ferment friendship. And knowing what morality requires is important knowledge. But is living a morally good life in itself—not under the heading of achievement, or as an instrument to pleasure or friendship, or in its connection with knowledge—an element of well-being?

This is not a question to which the answer seems to me directly apparent. The best I can do is approach the question indirectly, via what I call the sympathy test (I first proposed this test in Hooker, 1996, pp. 149–53).

Suppose we ask ourselves whether we are inclined to feel sympathy for someone whose life lacks a particular property. Sympathy is a judgement-sensitive attitude (on judgement-sensitive attitudes, see Scanlon, 1998, pp. 20-22). Our having sympathy for someone whose life lacks a particular property makes sense only if we judge that a life's lacking that property makes the life less beneficial to the person whose life it is than would be a life as similar as possible except that it has this property. So if we

2. An anonymous reviewer suggested that living a morally good life might not be much of an achievement if the agent were surrounded by people who always provided incentives for acting morally. That such incentives are always provided seems to me unrealistic.

do feel sympathy for someone whose life lacks the property, this attitude makes sense only if deep down we think that a life's having that property *is* an element of well-being. And if we do *not* feel sympathy for someone whose life lacks the property, one possible explanation is that deep down we think that a life's having that property is not an element of well-being.

We do feel sympathy for people whose lives lack pleasure, friendships, autonomy, significant achievement, or important knowledge without some sort of compensation in terms of a greater amount of one or more of these other goods. In contrast, we do not feel sympathy for people who fail to live morally good lives. One possible explanation of the absence of sympathy is that deep down we really do not believe that living a morally good life is an element of well-being.

Is the sympathy test a good one? If we do have sympathy for someone, we can legitimately make inferences about our beliefs. To be more specific, if we do feel sympathy for someone whose life lacks a particular property, then we must think that a life's lacking this property makes the life less beneficial to the person whose life it is than would be a life as similar as possible except that it has this property. The limitation of the sympathy test appears in cases where we do not feel sympathy for someone whose life lacks a particular property. Yes, one possible explanation for our lack of sympathy is that we really think that a life's having that property is not an element of well-being. However, another possible explanation is that something else prevents us from feeling sympathy. For example, we might think that the person under consideration deserves a life with lower well-being.

We might initially suspect that feeling sympathy for someone is difficult to combine with the condemnation and blame and indignation that we feel towards those we believe have failed to lead morally good lives. Moral blame is regularly accompanied by a kind of hostility, which can get in the way of sympathy.

And yet, blaming someone does not necessarily get in the way of feeling sympathy for that person. Sometimes we have to blame someone about whom we care very strongly. When we blame someone about whom we care very strongly, the blame *can* be accompanied by sympathy. For example, we might blame ourselves for something but at the same time feel sorry for the harm we have caused to ourselves.[3]

So far in this argument I have assumed that we do not feel sympathy for people who fail to live morally good lives. I have cast suspicion on the attempt to explain this lack of sympathy as an effect of blame for those who fail to live morally good lives.

3. I am grateful to Penelope Mackie for pointing this out to me.

Blame for people who fail to live morally good lives does not always prevent sympathy for them. So I *tentatively* surmise that the best explanation of our lack of sympathy for people who fail to lead morally good lives is that we do not really think that living such a life is a distinct element of well-being. (This conclusion should be tempered by the recognition that living a morally good life can be instrumentally beneficial to the person who lives it and can constitute a significant achievement.)

FOR WELL-BEING, FRIENDSHIP IS MORE THAN A KIND OF ACHIEVEMENT

Earlier, I argued that friendship should be listed as a distinct element of well-being. I have just now appealed to the sympathy test to argue that we do not really believe that living a morally good life is a distinct element of well-being, though living a morally good life is a kind of significant achievement, which is an element of well-being. With that conclusion in mind, someone might appeal to the fact that forming and sustaining friendships is also an important kind of achievement, albeit one less multi-dimensional than living a morally good life.

I do not see how it could plausibly be denied that that forming and sustaining friendships is also an important kind of achievement. Here is an example. Imagine someone named Frieda who has found the good in a friend named Markus and forgiven him and sustained her interest in him for decades. This really is an impressive achievement on her part, given how conflicted and moody and self-deluded and intermittently self-destructive Markus is. (In contrast, that he has continued to love her is no achievement on his part, given how breathtakingly easy she is to admire and appreciate.)

Now consider the following argument:

Premise 1: Forming and sustaining friendships is like living a morally good life in being an important kind of achievement.

Premise 2: Living a moral life is not a distinct element of well-being.

Premise 3: If forming and sustaining friendships is like living a morally good life in being an important kind of achievement, and if living a morally good life is not

a distinct element of well-being, then forming and sustaining friendships is not a distinct element of well-being.

Conclusion: Forming and sustaining friendships is not a distinct element of well-being.

Since this argument's conclusion does follow from the premises, we should assess the premises. Premise 1 is clearly true. Premise 2 is the conclusion we reached via application of the sympathy test, and so let us accept this premise. Premise 3 presumes that different things that are alike in one relevant respect are also alike in other relevant respects. We should not accept this premise, since different things are sometimes not alike in more than one relevant respect. Rejecting premise 3, we must reject the above argument as unsound.

To show that an argument is unsound is not to show that its conclusion is false. In the case of the above argument, however, I think there is an argument showing that its conclusion is false. In other words, this is an argument to show that forming and sustaining friendships is in fact a distinct element of well-being.

This argument begins with the premise that there is diminishing marginal benefit in each element of well-being (this argument was inspired by Hurka, 1993, pp. 84-96; 2011, pp. 166–74). For example, a life with no achievement but a lot of pleasure would benefit a great deal from gaining a significant achievement. Let us dub this achievement A of size S. Compare a very different life, one with lots of achievement already. This second life would not benefit as much from gaining the same achievement A of the same size S. In other words, the two lives we are comparing each gain an achievement A of size S; but, in the life where this is the *only* significant achievement, achievement A of size S constitutes a large benefit, and, in the life where there were already lots of other achievements, achievement A of size S constitutes a smaller benefit.

The diminishing marginal benefit to a life of its containing more instances of a kind of value of which it already contains a lot is relevant for the following reason. Imagine a life that already contained a lot of achievement but as yet no friendship. Suppose now this life gains one friendship. If friendship were not a distinct element of well-being but instead merely a subcategory of achievement, then a life that already contained a lot of achievement but as yet no friendship would *not benefit much* from the addition of one friendship. However, a life that already had a lot of achievement

but as yet no friendship would *benefit hugely* from the addition of one friendship. So friendship is not merely a subcategory of achievement but instead is a distinct element of well-being.

SEPARABLE-ELEMENT—VS—COMBINED-ELEMENT THEORIES

Let us turn from the question on what is on the list of elements of well-being to the question of whether the elements must be combined or need not be combined. We can agree that the very *best* life is one containing all the elements listed—pleasure, knowledge of important matters, friendship, significant achievement, and autonomy. Now we can distinguish between two kinds of objective list theory. These two kinds disagree about whether each of knowledge of important matters or friendship or significant achievement or autonomy on its own constitutes a benefit. Combined-Element Theories hold that only in combination do the listed items constitute benefits. Separable-Element Theories hold that the listed elements constitute benefits to the agent even when the agent does not have the other items. For example, on the Separable-Element Theory, knowledge of important matters (or significant achievement) constitutes a benefit to the agent even if it produces no pleasure for the agent.

Joseph Raz offers a particularly influential version of the Combined-Element view. Raz champions the partly objectivist view that well-being consists in the successful pursuit of worthwhile goals—in a slogan, 'subjective engagement with objective value'.

Raz's theory is partly objectivist because which goals are worthwhile is an objective evaluative question about values. '[G]oals are supported by approving judgement. ... [O]ne must regard one's goals in a way which ascribes to them desirability characteristics.' (Raz, 1986, p. 300) People 'engage in what they do because they believe it to be a valuable, worthwhile activity. ... To the extent that their valuation is misguided it affects the success of their life.' (Raz, 1986, p. 299) '[A] person's belief that his goal is valuable does not make it so.' (Raz, 1986, p. 344) People can think they are pursuing goals with desirability characteristics and yet be mistaken. People can unknowingly pursue valueless goals. And people who are successful in pursuing *valueless* goals do not thereby add to their well-being, according to Raz.

Raz's theory is not wholly objectivist. Facts about the individual subject's intentions and psychological engagement matter crucially. In Raz's view, someone would

fail to benefit from having good things in her life like knowledge and friendship if she did not have the right intentions and attitudes towards these things.

I agree with Raz that the kind of life that is *best* for any agent is one full of the successful pursuit of valuable goals. Furthermore, his ranking of active over passive states seems right. However, in contrast with Raz, I contend that pleasure can add to an agent's well-being even if this pleasure is not associated with the successful pursuit of worthwhile goals.

Suppose Ajay has a life with a given amount of successful pursuit of worthwhile goals. (For the purposes of my argument, it doesn't matter whether this amount is high, or low.) Now suppose that Ajay is given an increment of *passive* pleasure, not pleasure from the pursuit of worthwhile goals. Maybe he is introduced to a drink he can savour each night right before bed. Or perhaps he is blessed with particularly pleasurable dreams each night. In either case, hasn't his well-being increased, admittedly only a little, but still increased? I propose that, of any two individuals with equally successful pursuit of equally worthwhile goals, the one whose life contained more pleasure, even if this pleasure is only of a passive kind, has had greater well-being.

Raz anticipated this objection (1995, p. 7):

> *Not all pleasures contribute to one's well-being. ... I stretch myself on the beach and enjoy the warmth of the sun; I see a pretty rose, and enjoy the sight. My life is not better or more successful as a result. It is different if I am a beach bum, or ... a flower lover. But in that case the passive pleasure fits in with my activities, I am the sort of person who will make sure that there is room in his life for these pleasures. In that case the occasional pleasure contributes (if it does) to my well-being because it contributes (a tiny bit) towards the success of activities I am set upon. If I have no interest in sun on the beach or in flowers, these pleasures, while being real enough and while valuable as pleasures, do not contribute to my well-being. They have no bearing on my life as a whole. ... [T]he pleasure has a meaning in the life of the flower lover which differs from its meaning in the life of the one who is not. That difference makes it reasonable to regard the pleasure as active in one case and passive in the other. It is active where it meshes with one's general orientation in life.*

Of course I agree that pleasures matter more to some people than they do to others. And I agree that the pleasure of looking at a flower differs in meaning depending on whether that pleasure meshes with one's projects. But agreeing with Raz about

those matters does not preclude thinking that passive pleasures do add to well-being, even if less than active pleasures. Examples such as the one about Ajay suggest that passive pleasures do add to well-being. (That said, *perhaps* passive pleasures of the kind Raz is discounting are lexically inferior to goods of other kinds. What I mean is that no amount of these passive pleasures is large enough to be worth choosing over a significant amount of those other goods. For discussion, see Crisp, 1996, pp. 23–42.) Hence, I am willing to be a bit more subjectivist than Raz.

I am also willing to be a bit more objectivist than Raz. Suppose someone stumbles upon some important wisdom or 'falls into' a rewarding relationship. Even if obtaining important wisdom and having rewarding relationships were not this person's goals, do not the wisdom and the relationship this person obtains add to her well-being? Suppose now that she does not appreciate the value of that wisdom or relationship until very late in life. That the value was not appreciated earlier is sad. But when finally it is appreciated, the thought that our agent might have would be 'here has been this good in my life all along and I didn't realize it. Because of this good in my life, I was better off than I thought.'

Such thoughts are possibly correct. And when they are correct, their correctness does not depend on the agent's having them. Put crudely, it could be a true proposition that, because of a great unappreciated good pervading the person's life, this person had been better off than she realized. The truth of this proposition does not depend on the person's ever coming to believe the proposition. In other words, if the person dies without having ever even considered the proposition, the proposition might nevertheless be true. But if this proposition is true, then an agent's subjective appreciation of a good is not a necessary condition of that good's adding to her well-being.

Raz himself denies what he calls the 'transparency of intrinsic value', that is the thesis that 'a feature is intrinsically good only if, under normal conditions, the person (or other animal) for whom it is a good is content with its presence and prefers it to its absence'. (Raz, 1986, p. 269; cf. pp. 268, 321) I too deny the transparency of intrinsic value. Because I deny this transparency, I can hold that a person's well-being was greater because of some wisdom she happened upon or a relationship she fell into, although, throughout the time her life contained these elements, she was not content with their presence and did prefer their absence.

Again, Raz anticipated the idea. He wrote (1995, p. 6),

> *The condition of whole-hearted pursuit presupposes that even where the activity, commitment, or relationship is not one which the agent chose, or could abandon by choice, he is in control of the manner of his engagement in it. He has to direct his conduct in the light of his objectives and commitments, to guide himself towards his goal. ... In the main, the notions involved ... exclude resentment, pathological self-doubt, lack of self-esteem, self-hate, etc. One is acting whole-heartedly if one is not prey to one of these attitudes. Nothing else is required: no reflective endorsement of one's activity, no second-order desire to continue with it, etc. ... The fact that certain alternatives never cross one's mind may be a condition of having an appropriate attitude to people or activities.*

The concession Raz makes here is compelling. In order for an activity to benefit the person who engages in it, it is not necessary for the person to reflectively endorse the activity or to have a second-order desire to desire it. Where the agent's activity or relationship or other good is something he does not reject or resent, his successful engagement with it adds to his well-being.

Nevertheless, Raz's concession does not, I think, go far enough. His 'condition of whole-hearted pursuit' restricts the activities and relationships that constitute benefits to an agent to those activities and relationships that the agent does not doubt, reject, or resent. However, an agent's activity or relationship or some other good in his life could be something that he actively rejects or resents and yet this activity or relationship or some other good in his life could *still* constitute an addition to his well-being. Of course, the upset he experiences is a loss for him. But the loss should not blind us to the possible associated gain.

For illustration, imagine an agent who has a relationship that she did not choose but came to resent, and this resentment went on for years. But then finally she came to see that in fact this relationship was a blessing she hadn't recognized. Though not chosen, and long resented, this relationship, she discovers, turns out to be one of the most significant things in her life. Perhaps it turns out to be the only significant relationship in her life, maybe even the only significant good in her life. Neither the fact that she didn't choose the relationship nor the fact that she resented it for years precludes the relationship's constituting an important addition to her well-being.

Again, this is a truth that the agent might come to see. Then again, the agent might not be so fortunate as to come to see it. The agent's recognition of the truth

would be very good. However, this recognition isn't a necessary condition of the truth's obtaining.

I have contended that the limits of well-being extend farther both in a subjectivist direction and in an objectivist direction than Raz allows. I have argued that passive sensations (a subjective element) can be sufficient on their own to constitute at least a little well-being. I have also argued that objective goods such as knowledge and relationships, even if not chosen and even if resented, can on their own constitute an addition to well-being, though the distress they cause also constitutes a subtraction from well-being. Whether the addition is greater than the subtraction depends on the importance of the knowledge, the nature of the relationship, and the amount of distress. Perhaps it also depends on how much other knowledge the person has, and on how many other relationships her life contains.

To reiterate, the greatest well-being is constituted by whole-hearted, successful pursuit of worthwhile goals. What Raz has described—the combination of subjective and objective elements—is the ideal. But the separate elements constitute some benefit even when not combined with the others.

THE IMPORTANCE OF THE CONCEPT OF WELL-BEING

Once we analyze well-being into its elements, we encounter a challenge to the importance of the concept of well-being. The challenge is that, since anything important to say about well-being could be reduced to assertions about these elements, we can jettison the concept of well-being. [4] The challenge is to explain why the concept of well-being is not at best otiose. I will argue that it is not otiose.

Let us distinguish between intrapersonal trade-offs and interpersonal ones. In an intrapersonal trade-off, a loss in terms of a person's well-being occurs for the sake of a greater gain in terms of that same person's well-being. In an interpersonal trade-off, a loss in terms of one person's well-being occurs for the sake of a benefit to someone else.

In both intrapersonal cases and interpersonal trade-offs, there is difference between cases where gains and losses being traded off are in terms of a *single* element of well-being and cases where the gains and losses being traded off involve *multiple* elements of well-being. An example of an intrapersonal case involving a single element

4. This challenge was put to me at a conference on Happiness and Well-Being in Oxford in June 2013.

is my agreeing to take on some pain now for the sake of my having much less pain in the future than I would have if I did not take on the pain now. An example of an interpersonal case involving multiple elements is my taking some pain now for the sake of protecting your autonomy or your pursuit of significant achievements or your knowledge or your friendships.

In thinking about gains and losses involving multiple elements, we rightly find thinking in terms of well-being irresistible. The reason for this is, again, that significant achievement, important knowledge, friendship, autonomy, and pleasure have diminishing marginal value within a life.

Imagine that you could give your niece enough money for a very pleasant week at the beach. Alternatively, you could give your nephew enough money to enable him to delay for a week going back to his paid employment, and this week of free time would enable him to complete his musical score. Should you just weigh the amount of pleasure your niece would get against the magnitude of the achievement constituted by your nephew's completed score?

No, you should instead think about the size of the benefit your niece would get out of that quantity of pleasure and the size of the benefit your nephew would get out of his artistic achievement. And the sizes of the benefits to the different people of the pleasure and achievement would depend upon how much other pleasure your niece's life contains or how much other achievement is in your nephew's life. What matters really are the sizes of the benefits, not purely the quantities of pleasure and significant achievement taken on their own. In other words, what matters is the size of the additions to each's well-being, not the quantities of the elements added to each.

Acknowledgements: For helpful comments on an earlier version of this paper, I take pleasure in thanking Krister Bykfist, Roger Crisp, Tom Hurka, Eden Lin, Ingmar Persson, Julian Savulescu, Anthony Skelton, Valerie Tiberius, and two anonymous reviewers for this journal.

REFERENCES

Anscombe, E. 1958. *Intention*, Oxford, Blackwell Publishers

Crisp, R. 1997. *Mill on Utilitarianism*, London, Routledge.

———2006. *Reasons and the Good*, Oxford, Clarendon Press.

Feldman, F. 2004. *Pleasure and the Good Life,* Oxford, Clarendon Press.

Griffin, J. 1986. *Well-Being,* Oxford, Clarendon Press.

Hooker, B. 1996, 'Does Moral Virtue Constitute a Benefit to the Agent?' In R. Crisp, ed., *How Should One Live?*, Clarendon Press, 141–55.

Hurka, T. 1993. *Perfectionism*, Oxford, Clarendon Press.

———2011. *The Best Things in Life.* Oxford: Oxford University Press.

Nozick, R. 1974. *Anarchy, State and Utopia,* Oxford, Blackwell Publishers.

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

Quinn, W. 1993. *Morality and Action*, Cambridge, Cambridge University Press.

Rawls, J. 1971. *A Theory of Justice*, Cambridge, MA, Harvard University Press.

Raz, J. 1986. *The Morality of Freedom,* Oxford, Clarendon Press.

———1995. 'Duties of Well-Being', in Raz, *Ethics in the Public Domain*, Oxford, Clarendon Press.

Scanlon, T. 1998, *What We Owe To Each Other*, Cambridge, MA, Harvard University Press, 1998.

Sumner, L. W. 1996, W*elfare, Happiness, and Ethics*, Oxford: Clarendon Press.

Motives to Assist and Reasons to Assist: the Case of Global Poverty

SIMON KELLER

Victoria University of Wellington

ABSTRACT

The principle of assistance says that the global rich should help the global poor because they are able to do so, and at little cost. The principle of contribution says that the rich should help the poor because the rich are partly to blame for the plight of the poor. This paper explores the relationship between the two principles and offers support for one version of the principle of assistance. The principle of assistance is most plausible, the paper argues, when formulated so as to identify obligations that arise from the needs of particular identifiable members of the global poor, not from impersonal rules or values. Under that formulation, the principle can explain why knowledge of the circumstances faced by individual members of the global poor can have such a marked effect upon the willingness of the global rich to provide help, and can offer a better grounded motivational basis for helping the global poor. These are real advantages, the paper argues, and ones that cannot be matched by stories that focus upon the ways in which the global rich contribute to global poverty.

THE GLOBAL RICH AND THE GLOBAL POOR

The world's wealthier people are much, much richer than the world's poorer people. The size of the inequality is difficult to comprehend, but here are some credible estimates: 85% of the world's wealth is held by 10% of its population, the poorest 50% of the world's population holds 1% of its wealth, and the richest 20% of the world's population consumes 77% of its material goods (Shah 2011). Anyone reading

this paper is probably, in global terms, rich. "We" are the global rich and "they" are the global poor.

That the global poor are much poorer than us does not mean that they are poor absolutely, and it does not mean that they live bad lives. Many of the global poor, surely, live happier and more fulfilled lives than many of us. But many of the global poor *are* absolutely poor and *do* suffer profoundly as a result. Again, no brief description can be adequate, but here, again, are some credible estimates: 22,000 people die every day due to poverty-related causes, 1.4 billion people live on less than US$1.25 per day, 1.8 million children die each year due to diarrhea, and 400 million children lack access to safe water (Shah 2011). Poverty brings not only suffering but also vulnerability. Many of the global poor are unable to exercise meaningful control over the directions of their lives, and many are subject to economic and personal exploitation (O'Neill 2007).

The misfortune of the global poor is not unrelated to the fortune of the global rich, and it is plausible—though not uncontroversial—to say that the global rich are partly responsible for the plight of the global poor. Entrenched global poverty can be partly attributed, arguably, to an ongoing history of colonial, military, and financial interventions of rich countries into poor countries. We participate with the global poor in an integrated global economy whose structures and guiding institutions, arguably, are largely mandated by rich countries and ensure that global poverty is perpetuated (Pogge 2007, pp. 634-638; Singer 2009, pp. 30-35).

It is also plausible—though not, again, uncontroversial—to say that it is within our power to improve the lives of the global poor, and to do so at relatively small cost to ourselves. Compared with the amount of money spent by a person like you and me every day on luxuries, it does not cost very much to provide a poor person in a poor country with essential medicines, nutrition, and shelter, and there exist prominent charities that do a good job of transferring money from the global rich to the global poor (Singer 1972, p. 232; Singer 2009, pp. 15-19 and ch. 6). We may also help the global poor indirectly by donating money to political campaigns, pressuring our governments and corporations, and changing our habits as consumers, to give just a few examples (Singer 1972, p. 240; Singer 2009, ch. 7).

HELPING THE GLOBAL POOR: REASONS VERSUS MOTIVES

Supposing that the claims just mentioned are all correct—we are much richer than the global poor, many of them are badly off in absolute terms, their poverty is partly our fault, and we are in a position to help them—it is very plausible to think that we have a moral reason to help the global poor: that morally, all else equal, we *should* do something to help. But what is the nature of that moral reason, exactly? Here are two possible stories (Barry and Overland 2012, pp. 869-870; Barry and Overland 2013).

On one story, we should help the global poor because they are in need and we can help them, and at a relatively insignificant cost to ourselves. Our moral reason to help the global poor, runs the story, arises just from the facts about their plight and our ability to help alleviate it. This story rests upon a moral principle that we can call "the principle of assistance." The principle of assistance says that we should help the global poor because we can, and at little sacrifice (Singer 1972; Cullity 2004; Singer 2009).

On another story, it is a crucial additional fact that we bear blame for the plight of the global poor. Our moral reason to help the global poor, runs this second story, is that we owe them restitution for the harms that we—the global rich—have done them. The principle that lies behind this story can be called "the principle of contribution." The principle of contribution says that we should help the global poor because we have contributed to their poverty (Pogge 2002, pp. 11-20; Pogge 2007, pp. 633-634; Pogge 2010, pp. 20-15; Risse 2012, pp. 113-118).

Each story is plausible on the surface, and the principle of assistance and the principle of contribution are not incompatible. The two principles, as I have expressed them, are concerned only with moral reasons to relieve global poverty, but they each reflect something compelling about moral reasons to help others more generally. If you see a child who is in distress after falling off her bike, and if you can easily give her help, then you should help her—and you have extra reason to help her if it was you who caused her fall.

Many of us find it very natural, then, to think that we have strong moral reason, of one kind or another, to help poor people in other parts of the world. But those of us who believe that we have such a reason are unlikely to believe that we do a very good job of acting on it. Our governments do not do much, really, to relieve global poverty; the world's richest countries give around 0.29% of their gross national income in aid to impoverished countries (Organisation for Economic Co-operation

and Development, 2015). As individuals, many of us do nothing, and many of us do hardly anything. There are few who would accept the principle of assistance or the principle of contribution while also believing that the obligations of the global rich to the global poor are being fully discharged.

It is not just that we do not do enough to help the global poor. It is also that our patterns of helping the global poor do not appear to respond reliably to our reasons to help them. The global poor are always there and always in need of help, but we are more likely to be generous in response to identifiable, delimited individuals or events than in response to ongoing global poverty. We are more likely to give to a relief effort to help victims of a war, an earthquake, or a tsunami than to give to efforts to help people whose poverty is caused by long-term systemic factors. We are more likely to give if we take ourselves to be sponsoring a particular impoverished child, rather than to be one contributor to an effort to help many impoverished people. We are moved more by images and stories of particular poor people than by information about the overall problem of global poverty (Singer 2009, pp. 48-56). In all of these respects, our tendencies to help the global poor seem insufficiently responsive to the truly morally weighty features of the phenomenon of global poverty: its sheer size and its entrenchment. When it comes to helping the global poor, we may say, our reasons are not matched by our motives.

THIS PAPER

That is my description of the basic facts, including the basic moral facts, of the relationship between the global rich and the global poor. As mentioned, my description is open to dispute at several points, but still, it is plausible and widely believed and it sets the scene for this paper. My goal in this paper is to defend a version of the principle of assistance: one that I call the "individuals-based" principle of assistance, and that grounds our moral reasons for helping the global poor in considerations of particular individual members of the global poor. The principle of contribution, for all I say, may also yield moral reasons to help relieve global poverty, and other versions of the principle of assistance may do so too; the various stories are not, after all, incompatible. But I want to show that only the individuals-based principle of assistance can explain certain crucial facts about our motives to help the global poor, and that this is a reason to think the principle correct. If I am right, then the individuals-

based principle of assistance must be a big part of the truth about our moral reasons to help relieve global poverty.

My case for the individuals-based principle of assistance draws conclusions about *reasons* to help the global poor based on some claims about our *motives* for helping the global poor. It is one thing to talk about our reasons and another to talk about our motives. But, I want to show, these topics are linked. Different stories about our reasons yield different claims about what knowledge we have of our reasons, and our motives are likely to be different depending upon how much we know about our reasons. More specifically, I shall try to show that the individuals-based principle of assistance predicts that we in our present situation lack full knowledge of our moral reasons to help the global poor, and that it follows that our motives for helping the global poor are likely to be unreliable and sporadic, in much the ways we in fact find them to be. That is a reason for thinking that we take our reasons to be those posited by the individuals-based principle of assistance, and *that*, I shall go on to argue, is a reason for thinking that that principle is correct.

I should be clear about how I am using some terms. A "reason" to act, as I use the term, is a consideration that counts in favor of an action. By "counts in favor" I mean "*really* counts in favor"; on my way of talking, a "bad reason" is no reason at all. A "moral reason" is a *moral* consideration that counts in favor of a certain action. When I offer something as a moral reason for you to do something, I mean to explain why you morally should do it, all other things equal. That playing computer games all afternoon would be fun, for example, might be a reason to play computer games all afternoon, without being a moral reason.

A "motive" for action, as I use the term, is whatever moves a person towards performing an action. (Note that you might have a motive to perform an act but not in fact perform the act; you might try to perform the act and fail, or your motive might be outweighed by a different motive.) Sometimes, what moves you to perform an action might be a consideration that counts in favor of the action, and is hence also a reason to act: you might be motivated to brush your teeth by the thought that brushing your teeth helps avoid tooth decay, and the fact that brushing your teeth helps avoid tooth decay may be a consideration that in fact favors your brushing your teeth. Sometimes, your motive may be a consideration that does not in fact count as a reason: you might be moved to brush your teeth by the thought that brushing your teeth helps save the rainforests. And sometimes, your motive may involve no consid-

eration at all: you might be moved to brush your teeth by a bare urge or impulse or habit.

The paper's argument rests upon distinctions between different versions of the principle of assistance and the principle of contribution, and a conviction that making the distinctions is worthwhile. So I will start by saying how the principle of assistance and principle of contribution might be more precisely expressed, and why the choice between them, in their various forms, matters.

ASSISTANCE AND CONTRIBUTION

The principle of assistance says that we should help the global poor because they need our help and we can provide it, and at little sacrifice to ourselves. The principle of contribution says that we should help the global poor because we have contributed to their poverty. There are various strategies through which either principle could be connected with a wider claim about obligations to help others. It is an open question on the principle of assistance, for example, whether we have obligations to help the global poor even when doing so *does* require significant sacrifice, and it is an open question whether we have obligations to help anyone—poor or not—whom we could help with little sacrifice. Similarly, an advocate of the principle of contribution may say that we are obliged to give restitution for all harm to which we contribute, or only for harm to which we wrongly contribute, or whatever else.

Even as stated, though, the two principles posit reasons that apply in different cases and that sometimes pull in different directions. There may be people who live in severe poverty and whom we are able to help, but whose poverty cannot be blamed on the global rich; in such cases, the principle of assistance but not the principle of contribution posits reasons to help. There may be impoverished people whose poverty is not severe, by world standards, and whom we could only help with difficulty, but whose poverty is clearly our fault; in a case like this, the principle of contribution would appear to offer a reason to help, but the principle of assistance may not.

There may also be cases in which the two principles yield different claims about whom it is most important to help, or about whom we should help first. Depending upon exactly how they are expressed, the principle of assistance is likely to recommend that we be guided by the goal of giving as much help as we can to those who need it most, while the principle of contribution is likely to recommend that we be

guided by the goal of giving compensation to those to whom we have caused the greatest harm. These goals can conflict. Sometimes we may have to choose, for example, whether to help someone who is in greater poverty or someone for whose poverty we carry greater blame.

It matters for practical purposes whether we accept one of the principles and dismiss the other, and if we accept that there is truth in both principles, then it matters how we weigh them against each other. It also matters, in a number of respects, how exactly the two principles are formulated. Let me briefly mention two questions of formulation faced by the principles, before discussing at some length a third.

First, the principles need to say what counts as helping the global poor. The principle of assistance could tell us to make the global poor happier, for example, or to satisfy their basic material needs, or to make them less exploited, or to ensure that they have greater opportunities. The principle of contribution could tell us to compensate the poor with money, for example, or with whatever would best address their lack of happiness, or it could say that the global poor are now our responsibility in some broader sense.

Second, the principles need to say, at various points, whether they are speaking of groups or of individuals. To see how this makes a difference, imagine the following case.

Suppose that the philosophers at a university have an obligation to provide morning tea for the whole faculty. This will be an obligation held by the philosophers as a group. It does not matter which particular philosophers provide the morning tea, and it does not matter if some philosophers take no part in its provision. So long as the morning tea is provided by someone acting on behalf of the philosophers, the obligation is discharged. Our obligation to help the global poor could, like this one, be a group obligation. It could require that something be done on the part of the global rich, considered as a group, whether any particular member of the global rich contributes or not.

The philosophers at the university may also have an obligation to attend a staff meeting. This obligation will not be held by the philosophers considered as group, but will instead fall on each philosopher considered separately: each philosopher is obliged to attend the meeting. There are then really several obligations to attend the meeting—one for each philosopher—and those obligations cannot be discharged just through some philosophers acting in the name of all the philosophers. If our obligation to help the global poor is like the philosophers' obligation to attend a staff

meeting, then it is an obligation that falls separately upon each of us. To speak of the obligation of the global rich would then be to speak, really, of many obligations—one for each member of the global rich—and it would not be possible for all our obligations to the global poor to be discharged just through the acts of some rich people acting on behalf of all of us.

Just as an obligation can fall collectively upon a group or separately upon each member of a group, an obligation can be held towards a group or separately towards each member of a group. The obligation of the global rich could be an obligation to improve the lot of the global poor, considered as a class. Or it could be an obligation held towards each person within that class, so that there are really many obligations on the part of the global rich, one towards each member of the global poor.

The same goes for talk about the ability of the global rich to help the global poor. Are we talking about what could be done by the global rich, considered as a class, or about what could be done by the various individual rich people? Are we talking about an ability to help the global poor, construed as a group, or about various abilities to help various particular members of the global poor?

And, the same goes for talk about the contribution made by the global rich to the plight of the global poor. Are we talking about a contribution to the existence of a class of poor people, or are we talking about a contribution to the poverty of particular poor individuals? Are we talking about the contribution made by the global rich, considered collectively, or about the individual contributions of various individual rich people?

Depending upon how the principles answer these questions, they will tell different stories about where the obligations to help reduce global poverty fall, where they come from, and what it would take for them to be met. How the principles answer these questions also makes a difference to how plausible the principles are, and to how they can be compared and contrasted with each other—and also, I will say later, to how the reasons posited in the principles might be expressed in our motives.

UTILITARIANISM AND THE PRINCIPLE OF ASSISTANCE

I now want to spend some time discussing a third question faced by the principles. The question is about what kind of consideration grounds our moral reason to help the global poor, and whether that consideration is concerned with certain

individual people, or instead with certain values. I will start by explaining how the question arises for a utilitarian approach to the principle of assistance.

Utilitarianism, as I shall speak of it, says that the right act to perform under any circumstances is the one that produces the greatest possible balance of happiness over unhappiness. (This is one form of utilitarianism. There are others, but for the sake of illustration this one will do.) Utilitarianism, in conjunction with the facts about global poverty, leads quickly to a version of the principle of assistance. There are acts we can perform that will involve relatively small sacrifices in happiness for us while bringing about relatively large increases in happiness for the global poor, so by helping the global poor we can increase the overall balance of happiness over unhappiness. But consider two different stories about how you might come to utilitarianism in the first place.

A first story about how to get to utilitarianism begins with some claims about value. Happiness, you might think, is a good thing: it is good to be happy, and the world is better if it contains more happiness. And unhappiness, you might think, is a bad thing: it is bad to be unhappy, and the more unhappiness a world contains, the worse it is. You might think, furthermore, that where things apart from happiness and unhappiness are good or bad, their goodness and badness can be understood by reference to their contributions to happiness and unhappiness. Happiness and unhappiness, you might then conclude, are the only two things that have foundational value or disvalue, or that contribute in their own rights to making the world better and worse.

What will that tell you about rightness and wrongness? Well, you might accept the suggestion that people act rightly when they make the world better and wrongly when they make the world worse. If that is what you think, then you will go on to conclude that morally right acts are those that produce as much happiness and as little unhappiness as possible. And so you will have arrived at utilitarianism.

That story is condensed and stylized. The important point, for my purposes, is that it identifies a route to utilitarianism from a commitment about impersonal value. There is no mention within the story of the fact that the good and bad things—happiness and unhappiness—are always manifested in individuals. Only subsequent to the story, once we start working out what we actually need to do in order to maximize the balance of happiness over unhappiness, does it become relevant that we can only change facts about happiness and unhappiness by making a difference to the levels of happiness enjoyed by particular individuals. As far as the story goes, it is not im-

portant whether all the happiness and unhappiness is concentrated in a single individual, or how happiness and unhappiness are distributed across various individuals, or how many individuals there are. The concern in the first instance is with happiness and unhappiness, not with the individuals in whom happiness and unhappiness are to be found.

It is when we ground utilitarianism in this kind of story that we come most naturally to the thought that for utilitarianism, individuals are moral "buckets." For moral purposes, runs the thought, individuals matter only because they are receptacles in which happiness and unhappiness can be contained. Some expressions of the well-known objection that utilitarianism "does not take seriously the distinction between persons" are best understood as complaints about this utilitarian thought (Rawls 1972, p. 27). Utilitarianism misses something morally important, according to such complaints, because it addresses persons merely as contributors to impersonal value, not as separate individuals with their own lives to live.

Following this first story about how to argue towards utilitarianism, the principle of assistance will appear as a special case of a general injunction to distribute resources in ways that generate the most overall value. Our moral reason to help the global poor, on the story, is that by helping the global poor we can make it the case that the world contains more of the foundational good and less of the foundational bad. This is "values-based" utilitarianism, similar to what others call "teleological" utilitarianism (Kymlicka 2002, pp. 32-37). It supports a "values-based" principle of assistance. The reason revealed by the principle of assistance is a reason to respond properly to the value of happiness and the disvalue of unhappiness.

Now for the second story about how to argue towards utilitarianism. Instead of beginning with some claims about value, suppose that you begin by thinking about the individuals for which you have moral concern. To the extent to which you are morally motivated, which individuals seem to you to matter? You might answer by saying that the morally considerable individuals are all those who are sentient, or who have subjective experiences, or who have a point of view, or who can be better or worse off, or who can be happy or unhappy—where it is plausible to think that an individual who has one of these features has all of them. Humans and sentient nonhuman animals matter morally, you might think, but trees and bricks do not.

Thinking now about some particular morally considerable individual, ask what sort of treatment it merits. What is the right way to treat an individual that has its own point of view? There are many possible answers, but one of them is that the right

way to care about a morally considerable individual is to care about its happiness. Many individuals in the world merit moral concern, on this answer, and the right kind of moral concern, as addressed to any one of those individuals, is a desire that the individual enjoy happiness rather than suffering unhappiness.

For each individual, you may then think, we have a moral reason to do things that contribute to its happiness and detract from its unhappiness. Such reasons, you might conclude, are the most fundamental moral reasons there are.

Such reasons, however, conflict. We cannot always make everyone happy. Sometimes, the act that will maximize one person's happiness will set back the happiness of another; sometimes, we have no choice but to discriminate between the demands made by different morally considerable individuals. So, we need a rule that tells us what to do when we cannot do the right thing by everyone. There are many rules we could conceivably adopt, but one of them is this: when you cannot make every individual happy, do whatever produces the most happiness and least unhappiness overall. It is unfortunate that we sometimes need to sacrifice the happiness of some individuals to secure the happiness of others, but when it must be done—you might think—the way to do it is to see which individuals have what quantities of happiness at stake, and then to choose options yielding greater quantities of overall happiness over those yielding smaller quantities.

And that is the second story, again condensed and stylized, about how to arrive at a utilitarian criterion of right action. What matters for present purposes is that the story gets to utilitarianism from a concern with individuals. Unlike the first story about how to get to utilitarianism, this second story does not explicitly call upon judgments about impersonal value. Unlike the first story, this second story passes through a point in the argument at which it would be possible, having identified the morally considerable individuals, to say that we should be concerned not with their happiness, but instead with—say—their preferences, moral flourishing, or levels of autonomy; and it passes through a point at which it would be possible to reject the maximization principle in favor of a principle that tells us—say—to spread happiness equally, or to maximize the level of happiness enjoyed by the person who has least, or to refuse to make comparisons between the happiness levels of individuals at all.

Compared with the first story about how to argue towards utilitarianism, it is more difficult to accuse this second story of failing to take seriously the distinction between persons. The command to maximize the balance of happiness over unhappiness, on this second story, is offered as a way of coping with the fact that different

persons have different and competing points of view. It is never imagined that individuals are not separate, or that they matter only for their contributions to an overall impersonal value. Utilitarianism arrived at through this line of thinking can be called "individuals-based" utilitarianism; it has affinities with what others have described as utilitarianism based on "equal consideration of interests" (Kymlicka 2002, pp. 32-37).

Like the first story, this second story about how to get to utilitarianism yields a version of the principle of assistance, but the line of thought that takes us there is different. The poor people of the world are morally considerable individuals, just like us, so we ought to make them happy rather than unhappy. To contribute to their happiness, we need to make sacrifices in our own levels of happiness, and that is unfortunate. But, looking at the plights of those individuals, it is clear that the stakes for them are far greater than are the stakes for us. So, if we have a proper moral concern for those individuals, then we will see that we ought to make the relatively small sacrifices in our levels of happiness required to secure relatively large increases in theirs. Our reason to help the global poor is a reason to respond properly to the morally considerable individuals that make up the global poor. It is in this sense that we have arrived at an "individuals-based" principle of assistance.

TWO VERSIONS OF THE PRINCIPLE OF ASSISTANCE

The two stories about how to get to utilitarianism illustrate the distinction between the values-based and the individuals-based principle of assistance, but the distinction does not rest on anything special about utilitarianism.

Speaking generally, the values-based principle of assistance says that our moral reason to help the global poor is a reason to respond properly to certain values. Values-based utilitarianism, as I have imagined it, says that the values are happiness and unhappiness and the proper response is to maximize one and minimize the other, but other stories could be told. There can be other claims about which are the relevant values, and there can be other claims about what counts as a proper response to them.

An example of a non-utilitarian values-based principle of assistance can be found in Onora O'Neill's Kantian account of our obligations to the global poor (O'Neill 2007, see especially pp. 558-560). O'Neill endorses a version of the principle of assistance, saying that the fact that people are in poverty in itself provides us with an (imperfect) duty to do something to help them if we can. The consideration that

explains the existence of that duty is, for O'Neill, the value of autonomy: the plight of the global poor is morally significant because poverty undermines the preconditions of autonomy. The correct response to the value of autonomy, on O'Neill's view, is not to maximize it, but rather to live a life that honors it. For O'Neill, to honor the value of autonomy is, among other things, to recognize that humans need the support of others to achieve full autonomy, and hence to do something, here and there, to alleviate conditions that make autonomy more difficult to achieve.

On the individuals-based principle of assistance, our moral reason to help the global poor is a reason to respond properly to the various people who make up the global poor. The individuals-based version of utilitarianism, as I present it, says that the right response to those people is to try to increase their happiness, but, again, there are other stories to tell: other possible stories about what treatment people merit, and about how to make discriminations when not everyone can get the treatment they would ideally receive.

An example of a non-utilitarian individuals-based principle of assistance might be found in the view that our reasons to assist the global poor are generated by their natural positive rights. It might be held, for example, that every person has a right to have her basic needs met, if the resources required to meet those needs exist (Ashford 2007). The way to respond to any morally considerable individual, it might be held, is to respect her rights, and so for each person within the global poor we have a reason to do what we can to alleviate that person's poverty.

It is also possible to tell a Kantian story, different from O'Neill's, that generates an individuals-based principle of assistance. We could say that each person, as a rational being with the capacity for autonomy, demands in her own right whatever she needs in order to manifest genuine autonomy. That demand, we could say, then falls upon any person who could help make it the case that autonomy, for her, is achievable. As things are, many of the global poor face conditions that compromise their autonomy, and the global rich could alleviate those conditions at little sacrifice. So, we could conclude, our moral reason to help the global poor can be broken down into many different reasons to help: one for each of the rational individuals whose ability to live autonomously is compromised by her poverty.

TWO VERSIONS OF THE PRINCIPLE OF CONTRIBUTION

There is a distinction between the values-based and the individuals-based principle of assistance. We can make the same distinction between different versions of the principle of contribution, though doing so is more difficult and perhaps less natural. Here is how it could work.

As a first possible route towards the principle of contribution, you could begin with a concern for particular individuals within the global poor whose poverty is partly caused by actions of the global rich. Then you could make the judgment that each of those particular individuals is owed restitution for the harms she has suffered, and hence that the global rich should help the global poor. The basic moral relationship behind the principle of contribution, so understood, is between the global rich, on one hand, and many members of the global poor considered separately, on the other. For each impoverished person whose poverty is partly caused by the actions of the global rich, there is a moral reason for the global rich to help that person. So this approach to the principle of contribution is individuals-based. Our moral reason to help the global poor, on this approach, really comes down to our many moral relationships with the many members of the global poor whom we have harmed, and to the many reasons that emerge from those many relationships.

As a second possibility, you could begin with the thought that the global rich, through their actions toward the global poor, have failed to respond properly to certain values. You might say, for example, that happiness is valuable and the global rich have, through their treatment of the global poor, spread unhappiness; or you might say that autonomy is valuable and the global rich have, through their treatment of the global poor, undermined autonomy. Where someone has failed to respond properly to a certain value, you might add, she can become obliged to make amends by doing what she can to restore the value where it was damaged. We might show proper respect for the value of happiness, for example, by trying to restore the happiness of the people whom our past actions have made less happy; and we might show proper respect for the value of autonomy by helping those whose autonomy we have damaged in the past to enjoy greater autonomy in the future. You might go on to say that we have contributed to the unhappiness and undermined the autonomy of the global poor, and so we now have a special responsibility to help the global poor, as an expression of respect for the values of happiness and autonomy. If that is how

you explain our moral reasons to help the global poor, then you have a values-based version of the principle of contribution.

I hope that the distinction between the values-based and individuals-based principles is now reasonably clear. It may become clearer as I explain the use to which I want to put it. First, though, I want to turn to a different topic: our knowledge of reasons, and how that knowledge is reflected in our motives.

KNOWLEDGE OF REASONS

It is possible to know that you have a reason to do something without knowing what that reason is. To put it another way, it is possible to know about the existence of a reason while lacking knowlege about its nature (Keller 2013, pp. 87-90).

Suppose you get some medicine from the doctor, and she tells you that you must take it with food. If I ask you, "Should you take the medicine with food?" then you may answer, "Oh yes: the doctor was very clear about that, and she's always right. I should definitely take it with food." If I then ask, "Why must you take the medicine with food?" then you may answer, "I have no idea. The doctor didn't tell me." You know about one reason to take the medicine with food—the fact that the doctor told you to—but the question "Why must you take the medicine with food?" is, in this context, a request for a different kind of consideration: a medical consideration, the one that lies behind the doctor's instruction. When it comes to that reason, it is natural to say that you know you have the reason, but you do not know what it is. You know that you should take the medicine with food, but you do not know why.

Here is a more involved example. Imagine that you and I are at a crowded market. We are good friends, and you trust me to tell you the truth. We are separated, we both have cell phones, and we are avid texters—imagine. You receive a text from me, saying "Come and meet me at the top end of the market right away." You reply, "OK. Why?" I reply, "Just come quickly! You have a very strong reason to come to the top of the market right away!"

At this point, assuming that I am texting truly and sincerely and you are justified in trusting me, you know you have a strong reason to go to the top of the market as fast as you can: a reason, again, beyond the fact that you have just been told to go there by someone you trust. Yet, you do not know what the reason is. You do not even know what kind of reason it is. It could be a moral reason; perhaps someone is in danger and I need your help to save her. It could be a reason of self-interest; perhaps

someone is giving away gifts at the top of the market, or perhaps I have seen a troublesome ex-partner of yours at the market and I am trying to help you avoid him. Or perhaps there is a beautiful sunset you will see from the top of the market, or perhaps someone you would like to meet is at the top of the market, or perhaps I am feeling panicky at the top of the market and need your immediate comfort. Any of these considerations would count as a reason for you to come to the top of the market, but they are all different kinds of reasons, and you cannot discriminate between them. All you know is that you should hurry to the top of the market, for some reason or other.

FROM KNOWLEDGE OF REASONS TO FULLY SHAPED MOTIVES

Think about the motives you will feel as you hurry through the market after receiving my text. Knowing that you have a good reason to get to the top of the market, you might move as fast as you can. But you will not find any determinate sense of urgency in your motives, because you will not know why you are acting. You are likely to feel a little bemused, alienated, or detached. Your rushing to the top of the market will not quite make sense to you. You will have a motive, but your ignorance about the nature of your reason will be reflected in your motive. As I will put it, your motive will not be "fully shaped."

Suppose that as you walk towards the top of the market, I continue to send you text messages. I tell you, "It's in your best interests to come to the top of the market straightaway!" Then, "There's a record for sale that you'll want to buy before anyone else gets their hands on it." Then, "It's by a seventies North American folk revival artist"; then, "It's a Gordon Lightfoot record"; then, "It's Sundown"; then, "It's an original mint condition US-issued LP"; then, "It's only $20"; then, "It has a yellow price sticker"; and finally, "The price sticker is on the top left-hand corner."

As you receive all this information, one piece at a time, you will gradually learn more about your reason to come to the top of the market as fast as you can. And as you learn more about your reason, your motive will become sharper. By the time you have received the last of my text messages, you will be fully engaged with your act of walking quickly towards the top of the market. You will fully understand what you are doing and why you are doing it. You will no longer be grasping for further information, needed to make your actions make sense. Your motive for walking to the top of the market will be fully shaped.

As a motive comes closer to being fully shaped, in the sense in which I am interested, it does not necessarily get stronger. As you learn more about why you should hurry to the top of the market, your motivation may in fact get weaker, as various more urgent possibilities are excluded. Once you know that you will be buying a record, not helping me cope with a panic attack or saving a stranger from danger, you may judge that you have a weaker reason than you had suspected, and this may be reflected in your having a weaker motive.

To have a fully shaped motive, rather, is to have a motive whose quality is not open to change with further information about your circumstances—meaning information that fills out the details of your present situation, not information that reveals that the situation has changed. Your motive for hurrying to the top of the market is fully shaped once you have all the information you need in order to have a clear and complete sense of why your hurrying to the top of the market makes sense; it is fully shaped once further information becomes irrelevant. So, for example, it may be that you want to buy a copy of Sundown, and once you hear that a copy is for sale, that is enough for you to be fully motivated to go buy it. The record's time and place of issue, its price, and the color of its price sticker may be matters of indifference to you, so that learning about them makes no difference to your motive to get to the stall and buy the record.

(Further information may yet make a difference to your motives not by changing the quality of this motive, but instead by adding another, different motive. If you learn that the person working at the record store is a good friend of yours, trying to make a start in the used music business, then you may come to be even more strongly motivated to come to the stall and buy the record, as a way of supporting him in his new venture—but now because you have two different reasons, and two corresponding motives, not just one.)

We can learn something about what you take to be your reason for performing a certain action by seeing how your motive for performing the action varies with new information. If you need a certain piece of information before you can have a fully shaped motive, then that piece of information would seem to play a role in your reason for acting, as you see it. If a certain piece of information would make no difference to the quality of your motive, then it would seem not to be relevant to the reason you take yourself to have. One way to see what you take as your reason to act is to ask what information you need in order for your motive to act to be fully shaped.

Returning to the example: if you come to have a fully shaped motive once

you hear that the record being sold at the top of the market is Gordon Lightfoot's Sundown—if the information beyond that makes no difference to the quality of your motive—then that suggests that you take it as your reason to rush to the top of the market that by doing so you will get a copy of Sundown. If, instead, you have a fully shaped motive once you hear that someone is selling a record by a North American folk revival artist—if it makes no difference to your motive to hear more about which artist or album it is—then that would suggest that your reason for rushing to the top of the market, as far as you are concerned, is just that you have the chance to buy a record by some North American folk revival artist or other. Alternatively, your motive might be fully shaped only once you hear that the record available is an original mint condition US-issued edition of Sundown, available for $20; that would suggest that you take it as part of your reason that you can get this edition of the album at this cheap price. Or, conceivably, you may have made a bet with someone that you could find a copy of Sundown with a yellow price sticker in its top left corner; in that case, perhaps, you will take it as an essential part of your reason to buy the record that it has a yellow price sticker in its top left corner, and so you will not have a fully shaped motive for buying it until the news about the sticker comes through.

KNOWLEDGE OF REASONS TO ASSIST

Back, now, to moral reasons to help the global poor. What do we know about the plight of the global poor and our ability to help? Most of us are aware that there are many impoverished people in the world. We are aware that many people suffer hunger, disease, and exploitation, and that many people lack freedom and autonomy and cannot fulfill their most basic needs. We also know that we are able to make those people better off. We know, for example, that if we donate some money to Oxfam, then our money will be used to improve the lives of some people among the global poor.

Most of us do not know anything much, however, about the particular people who make up the global poor. Most of us are acquainted with only a few members of the global poor at most. Regarding the vast majority of individuals among the global poor, we do not know who they are, or exactly what plights they suffer, or exactly what they most need for their lives to be improved. When we give money to Oxfam we know it will be used to help some people, but we cannot name names. As a result, we are not equipped to feel any particularized concern for the individuals among the

global poor. We are not able to have thoughts about particular poor people; when we think about helping the global poor, we are not able to think about anyone whom we are helping.

What do we then know about our reasons to help the global poor? The values-based principle of assistance and the individuals-based principle of assistance give different answers.

On the values-based principle of assistance, we would appear to know everything we need to know in order to have full knowledge of our moral reasons to help the global poor. We know that many people live in poverty, we know that the preconditions of autonomy are not met, we know that basic needs are not being fulfilled, we know that there is unhappiness and exploitation—we know all about the values at stake. We know, for example, that suffering is bad and that a donation to Oxfam will reduce suffering; if our reason to donate is a reason to reduce suffering, then in our present state of knowledge it is fully revealed.

On the individuals-based principle of assistance, however, we would appear to be in one of those situations in which we know we have a moral reason to act but do not know what the reason is. We know we have moral reasons to help the global poor. We know that there are many individuals in the world who demand our assistance, and whom we are able to assist. We know that for each of these individuals, we have a moral reason to help—but we do not know who those people are. We are not able to identify the individuals whose plights demand our assistance, so we are not able fully to identify our moral reasons to help them.

In light of the considerations discussed earlier, it follows that the values-based and individuals-based principles yield different predictions about our motives to help the global poor. On the values-based version, we can have fully shaped motives for helping the global poor. On the individuals-based version, we cannot. In being moved to help the global poor, on the individuals-based principle, we will lack a determinate sense of why our acts make sense. Particular individuals among the global poor are the sources of our moral reasons, but our motives do not have particular individuals on which to focus.

As mentioned earlier, most of us, many of us think, do not do as much as we should to help the global poor. There are many explanations of why this should be. We are selfish; we are weak-willed; we have other priorities; we feel helpless in the face of the problem of global poverty; we are ignorant, often willfully; we are tempted by claims that any contribution we might make will only make things worse; and so

on (Singer 2009, ch. 4). There are also many explanations of why when we do help the global poor, we are more likely to be moved by particulars than by generalities. The general facts about global poverty are too big and complicated for us to comprehend; our emotions are engaged by personal stories, not by statistics; it is easier to forget a general fact than a personal story; personal stories reveal the suffering caused by global poverty in ways that general facts cannot; and so on. But the individuals-based principle of assistance, placed alongside some broad truths about how our knowledge of our reasons is reflected in our motives, offers a further explanation.

If the individuals-based principle of assistance is correct, then we need to know about particulars, not just generalities, in order to know exactly what moral reasons we have to help the global poor, as opposed to knowing merely that the reasons exist. Lacking such knowledge, it is unsurprising, in light of broader facts about motives and knowledge of reasons, that we should fail to do what we know we have reason to do. As we learn more about particular individuals among the global poor and the struggles they face, and as we imagine ourselves to be able to help particular identified impoverished people, it is predictable that we should feel more urgently gripped by our reasons to help and more likely to act on them.

It is not surprising, if the individuals-based principle of assistance is correct, that when we are faced only by the general facts about global poverty, we feel listless and unmotivated. We can see that moral reasons to help are there, but we feel removed from them, because we do not know exactly what they are. But we come into closer contact with those reasons, or at least feel that we do, when we are presented with more details about the plights of the global poor and of our ability to help: when we know about a particular disaster and the challenges it raises, when we are offered the chance to sponsor a particular child or development project, and when we are presented with detailed and personalized stories. So under such circumstances, goes the explanation, we are more likely to do something to help.

The explanation is not incompatible with other explanations of our habits of helping the global poor. It in fact fits nicely alongside them. Perhaps our emotions are more vividly engaged by particulars because we respond more vividly to our reasons the more we know about their nature. Perhaps we fall into weakness of will partly because our reasons to help the global poor are too distant to engage our motivations in the face of other temptations.

IN DEFENSE OF THE INDIVIDUALS-BASED PRINCIPLE OF ASSISTANCE: PART 1

In offering this explanation of our motives to help the global poor, as so far expressed, we do not need to suppose that the individuals-based principle of assistance is correct. We need only suppose that most of us at some level take it to be correct. Our motives are informed not directly by our knowledge of our true reasons, but instead by our knowledge of the considerations that we take to be our reasons. (Even if you do not really have a reason to buy the copy of Sundown at the market, if you think you have such a reason, then that is enough to explain the quality of your motives as you receive the various pieces of information about its availability.) To put it another way, the explanation offered so far constitutes evidence that we take our moral reasons to help the global poor to be grounded in considerations of individuals, not of values—but still, we could be wrong.

That said, there is a good case for thinking that sometimes, at least, if we can show that we usually take a certain consideration as a reason, then we have strong evidence that it is in fact a reason. First, there is a case for being conservative, in a certain respect, in drawing conclusions about our reasons to help the global poor. Once we have accepted that we have moral reasons—of one kind of another—for helping the global poor, and once we have seen that we treat certain considerations, rather than others, as our reasons, and when we do not have any cause to think that in taking those considerations as reasons we are making a mistake: well, at that point in the argument, it is most reasonable to be conservative, and to presume that the reasons we take ourselves to have are the reasons we really have.

Second, there are times at which we seem, in having certain motives, to be getting things right. To have a moral motive, sometimes, is to treat a given consideration as morally powerful. When you see a child who has fallen off her bike, and you find yourself motivated to go and help her, you do not experience your motive merely as an urge, but also as a perception of what would be an appropriate response to the situation. Where a moral motive appears to be correct, or appropriate, or called for, it has some authority. If it can be shown that to have that motive is to take a certain consideration as a reason, then we respect the authority of the motive by accepting, in the absence of any evidence to the contrary, that it is not deluded: that the consideration treated as a reason really is a reason (Keller 2013, pp. 25-27, 84-87).

Coming to have a motive to help a member of the global poor upon learning

about the details of her circumstances or to help a group of people upon learning that they are affected by a disaster is plausibly an experience that carries moral authority, in the relevant respect. It is not an experience of being emotionally manipulated or of coming to have a brute or alien urge. It is rather an experience of seeing more clearly why you should do more to help this person, or to help the people affected by this disaster. The motives generated appear to be appropriate and to count as proper responses to the situations. The motives carry authority, and so by learning more about the reasons that you posit in having those motives, we can learn more about what reasons are really there.

It seems to me, then, that the story told by the individuals-based principle of assistance—the story on which our motives to help the global poor are explained as responses to varying levels of knowledge about reasons grounded in considerations of individuals—should be accepted as a story about what reasons there are to help the global poor, not just about what reasons we think there are. The individuals-based principle of assistance then has a virtue that the values-based principle of assistance lacks. It is when we learn more about individuals, not when we learn more about values, that we learn more about our reasons for helping the global poor.

The explanation offered by the individuals-based principle of assistance is not a straightforwardly empirical explanation. It seeks to explain observed empirical phenomena, but it does so by employing a substantive normative claim (about the nature of reasons) plus a moral psychological claim (about the impact of our knowledge of reasons upon our motives). That said, the explanation is also not a *moral* explanation, if by that is meant a justification or excuse. When you know that you have a strong moral reason to perform a certain act, you *should* perform that act, even if you do not know what the reason is. When you know that there are many impoverished people who morally demand your assistance, you *should* render assistance, even if you do not know who those people are. The explanation does not show that we act rightly when we fail to act upon our reasons to assist the global poor. Instead, it places that failure within a broader, independently understandable phenomenon—one that has nothing specifically to do with global poverty, or indeed with morality.

IN DEFENSE OF THE INDIVIDUALS-BASED PRINCIPLE OF ASSISTANCE: PART 2

If what I have said so far is correct, then the individuals-based principle of assistance tells an appealing story about our knowledge of our reasons to help the global poor and how that knowledge is reflected in our motives. The individuals-based principle of assistance is able to tell that story because it grounds our reasons to help the global poor in considerations of individuals, not of values. The principle of contribution, as mentioned earlier, is most naturally expressed as an individuals-based principle. On the individuals-based principle of contribution, our reasons to help the global poor are grounded in considerations of particular impoverished individuals to whose poverty the global rich have contributed. It may appear, then, that the individuals-based principle of contribution can offer a story just as attractive as the story told on the individuals-based principle of assistance.

For the individuals-based principle of contribution, the story would go like this. We can know that we have reasons to help the global poor just by knowing that many people are greatly impoverished and that we, the global rich, contribute to their poverty. Knowing that we have such reasons, however, is not the same as knowing what those reasons are. To have full knowledge of any of our reasons we need to know about the circumstances of particular individuals among the global poor, and we need to see how the poverty of those individuals, in particular, results from the acts of the global rich. Learning more about particulars of global poverty, and especially about our contribution to it, can then, on the principle of contribution, take us closer to full knowledge of our reasons to help the global poor. And that in turn can explain why learning more about particulars leads to our motives being more fully engaged, and why we are ultimately more likely to be moved by particulars than by generalities. The individuals-based principle of contribution, if this story succeeds, can do just as well as the individual-based principle of assistance.

The story told on the individuals-based principle of contribution is different in a small but significant respect from the story told on the individuals-based principle of assistance. On the individuals-based principle of assistance, you can know all about your reason to help a member of the global poor just by knowing about her poverty and your ability to help her. On the individuals-based principle of contribution, you also need to know that her poverty is partly due to the acts of the global rich. It is of course possible to learn that a person has suffered at the hands of the global

rich, and possession of that information can add to the sense that something really must be done to help that person. That a person's poverty is caused by the acts of rich people can make her poverty all the more morally outrageous. Learning that the global rich have contributed to a person's poverty, however, is different from learning that *you* have contributed to her poverty—even if you are a member of the global rich.

You may find that a particular person's poverty results from the acts of the global rich. It is usually not possible, however, to find that her poverty results from any acts performed by you. You might find that the person's poverty is caused partly by the acts of certain rich governments and corporations and by global political forces that lead back to the global rich, but that is not to say that anything you have done has made any difference at all to her circumstances. That is not how these large global causal processes work. The global rich caused her poverty, and you are a member of the global rich, but it does not follow that you caused her poverty. Anything you have done, probably, is irrelevant.

As a result, there is a considerable distance between finding, following the principle of contribution, that the global rich ought to do something for an individual and finding that you ought to do something for that individual. You are one member of the global rich, but you cannot discharge an obligation on behalf of the entire global rich: not when the obligation is accrued in response to large-scale collective actions and forces. And it is possible for the global rich to discharge its obligation without you doing anything at all. Learning that someone's poverty is due to actions of the global rich might lead you to believe that the global rich have an obligation to provide restitution, but that is not to believe that her poverty is due to your actions or that you have any resultant obligation.

A similar relationship holds in the opposite direction. As a member of the global rich, you may be able to identify actions of yours that harm the global poor. Your habits of consumption, your use of energy, your support of various policies, governments, and corporations—all of these may count as acts that help reinforce global poverty. But to say that you harm the global poor is not to say that you harm any particular identifiable member of the global poor, and, probably, it is not possible even in principle to find any particular person who is worse off because of what you in particular have done. That, again, is not how the contribution of the global rich to the poverty of the global poor works.

None of this is to show that the principle of contribution is mistaken, or that we do not accrue obligations as a result of our individual contributions to acts that

harm the global poor. The point is just that the story told on the principle of contribution cannot generate a direct moral relationship between a particular individual among the global poor and a particular individual among the global rich. Learning more about the circumstances of particular members of the global poor cannot pinpoint you as a person who has resultant reasons to help the global poor, arising from the contribution to global poverty made by the global rich. So when it comes to the phenomenon at issue—the power of knowledge of particulars in forming our patterns of motivation to help the global poor—the individuals-based principle of contribution cannot adopt the explanation offered on behalf of the individuals-based principle of assistance. Considerations of our motives in responding to global poverty ground an argument for the individuals-based principle of assistance, but not for the individuals-based principle of contribution.

CONCLUSION

I have tried to articulate and make plausible a principle that explains why we, as members of the global rich, have reason to help the global poor. The principle is the individuals-based principle of assistance: it says that we should help the global poor because we can, and at little sacrifice, and it grounds this claim in considerations of the particular individual persons among the global poor who would benefit from our help. It depicts our moral reasons to help the global poor as reasons that fall upon each of us considered separately, and that take individual members of the global poor separately as their objects.

My case for the principle is that it yields a plausible story about what we know about our reasons to help the global poor and about how our knowledge of our reasons is reflected in our motives. You are most likely to be moved to help a person in need when you are vividly aware of the person's needs and of a course of action by which you could help ensure that those needs are met. On the story I tell, this fact about our motivations is not merely psychological, but emerges predictably from the structure of our reasons and our ability to have knowledge of our reasons. Placed alongside some general truths about how our motivations are engaged by knowledge about our reasons, the individuals-based principle of assistance can explain why our motivations to help the global poor are most pointed and powerful when we are confronted with facts about particular people and particular circumstances.

The individuals-based principle of assistance can be distinguished from the

values-based principle of assistance and from the principle of contribution, and it is the only one of those principles, I have tried to show, that can tell the plausible story identified. It follows that the individuals-based principle of assistance gives the best explanation of the reasons held by individuals among the global rich to give help to the global poor. To that extent, there is good reason to accept the individuals-based principle of assistance as correct and indispensable.

It does not follow, however, that alternative principles are incorrect. There are other questions to be asked about reasons to help the global poor, apart from the question of what we as individuals among the global rich should be doing. Among them are the questions of what reasons the governments of rich countries have to help the global poor, and what corporations and other rich institutions should do, and how the arrangement of global institutions should be changed in response to the plight of the global poor. These are questions about what should be done by collective agents and institutions. Regarding these questions, perhaps, the principle of contribution is more plausible, as it can focus upon the contribution to global poverty made by larger entities and forces, and upon the global poor considered as a class.

Let me emphasize in closing that the story I have told about our reasons to help the global poor is not intended to show that our reasons are not as strong as we might have thought, or that they only become powerful when we happen to be faced by knowledge of particulars. When we confront our obligations to the global poor, we find ourselves in an epistemic situation that involves a certain kind of motivational barrier. Most of us, most of the time, know that we have reasons to assist the global poor, but we do not know what those reasons are, exactly; and it is a general truth that mere knowledge of the existence of a reason is not as motivationally gripping as is knowledge of the nature of that reason. That barrier appears to be a real and consequential barrier—by positing the existence of the barrier, we can explain something about how and when people come to be motivated to do things that help the global poor—and in recognizing it, perhaps we can come better to see why it should and could be overcome.

Acknowledgements: I am grateful for helpful comments from Nick Agar, Christian Barry, Garrett Cullity, Ramon Das, Ole Koksvik, Gerhard Overland, Vanessa Schouten, Dan Weijers, and two anonymous referees.

REFERENCES

ASHFORD, E., 2007. The Duties Imposed by the Human Right to Basic Necessities. In: T. Pogge, ed. *Freedom from Poverty as a Human Right*. Oxford: Oxford University Press, 183-218.

BARRY, C. and OVERLAND, G., 2012. Are Trade Subsidies and Tariffs Killing the Global Poor? *Social Research* 79 (4), 865-896.

——— 2013. How Much for the Child? *Ethical Theory and Moral Practice*, 16 (1), 189-204.

CULLITY, G., 2004. *The Moral Demands of Affluence*. Oxford: Oxford University Press.

KELLER, S., 2013. *Partiality*. Princeton: Princeton University Press.

KYMLICKA, W., 2002. *Contemporary Political Philosophy* 2nd ed. Oxford: Oxford University Press.

O'NEILL, O., 2007. Kantian Approaches to Some Famine Problems. In: R. Shafer-Landau, ed. *Ethical Theory: An Anthology*. Malden, MA: Blackwell, 553-564.

POGGE, T., 2002. *World Poverty and Human Rights*. Cambridge: Polity.

———2007. Eradicating Systematic Poverty: Brief for a Global Resources Dividend. In: H. LaFollette, ed. *Ethics in Practice* 3rd ed. Malden, MA: Blackwell, 633-646.

———2010. *Politics as Usual*. Cambridge: Polity.

RAWLS, J., 1972. *A Theory of Justice*. Oxford: Clarendon Press.

RISSE, M., 2012. *Global Political Philosophy*. Basingstoke: Palgrave Macmillan.

SHAH, A., 2011, *Poverty Around the World* [online]. Available from: http://www.globalissues.org/article/4/poverty-around-the-world#WorldBanksPovertyEstimatesRevised [Accessed 8 June 2015].

ORGANISATION FOR ECONOMIC CO-OPERATION AND DEVELOPMENT, 2015. *Development aid stable in 2014 but flows to poorest countries still falling* [online]. Available from: http://www.oecd.org/dac/stats/development-aid-stable-in-2014-but-flows-to-poorest-countries-still-falling.htm [Accessed 8 June 2015].

SINGER, P., 1972. Famine, Affluence and Morality. *Philosophy and Public Affairs* 1 (3), 229-243.

——— 2009. *The Life You Can Save*. Melbourne: Text Publishing.