PERSON-NEUTRALITY AND THE SEPARATENESS OF PERSONS

Matt Zwolinski

1.

People make trade-offs in their lives every day. Dieters give up tasty foods for the sake of better future health, entrepreneurs risk present dollars and energy for the sake of greater future returns, and so on. Our trade-offs are not always rational, but so long as it is our own time and energy we are wasting, or our own health we are risking, there is, in most cases, nothing immoral about the trade-off. Our resources are, for the most part, ours to do with as we please.

Not all trade-offs are like this. What makes the trade-offs described above acceptable from a moral point of view, irrational though they may be, is that the costs and benefits of the trade-off occur within a single life. An entrepreneur who sought, without your consent, to use your money for her financial gain would be acting outside her rights, even if her investment was guaranteed to be profitable. From a moral perspective, the size of her potential benefits vis-à-vis your cost does not matter. Your costs cannot be compensated by her benefits.

Like most things in philosophy, these claims are not uncontroversial. Some people think the size of the benefits does matter. Utilitarians, in fact, seem to think the size of the (net) benefit is *all* that matters. For the utilitarian, there really is not that much of a difference between *intra*personal trade-offs and *inter*personal trade-offs. In either case, our main concern should be simply to do as much good as possible. The way that good is distributed, either among the various time-slices of one's life, or among the lives of separate individuals, is of no fundamental importance.

Most non-utilitarians think there is something wrong with this conclusion. It is the result, some have said, of supposing "that mankind is a super-person, whose greatest satisfaction is the objective of moral action,"[1] or of supposing that there is some sort of "social entity,"[2] whose good is augmented by the sacrifices of society's individual members. In short, it is the result of failing to take seriously the fact that society is composed of distinct individuals, each with their *own* plans, ambitions, and claims against others. It is a result of failing to see that there is a difference between what is rational within a life and what is moral across lives. To blur the distinction between intrapersonal and interpersonal sacrifice is, in the eyes

of utilitarianism's critics, to ignore the separateness of persons.

This criticism has been enormously influential. But what exactly does it mean? And what lessons does it hold regarding the proper role of trade-offs in a theory of distributive justice? In this paper, I want to argue that appeals to the separateness of persons—at least as they are typically stated—cannot stand up to the normative burden that has been placed on them. In order to be successful, appeals to the separateness of persons would need both to point out some feature of utilitarianism that is responsible for its intuitively objectionable positions on issues of distributive justice, and to show how that feature somehow involves a failure to take seriously the separateness of persons. I do not believe that any of the currently popular appeals to separateness have accomplished both of these goals simultaneously. As they are currently stated, then, appeals to the separateness of persons do not constitute any real argument against utilitarian theories. They are, at best, promissory notes for future argument.

2.

What does it mean to claim that a moral theory ignores the separateness of persons? The natural interpretation would seem to be that there is some metaphysical *fact* of separateness that utilitarians have overlooked. People are separate in such-and-such ways, and utilitarians have ignored this fact. But this interpretation is surely far too uncharitable to ascribe—either to utilitarians or to their opponents who appeal to the separateness of persons as an argument against utilitarianism. The fact that there is more than one person in the world is just too obvious to suppose that utilitarians have been unaware of it, as is the fact that each of these persons has their own distinct set of ends, interests, or whatever. If we are to take appeals to the separateness of persons seriously, we must not read them as claiming that utilitarians have ignored some fact of this sort.

But neither should we understand appeals to the separateness of persons as being purely normative claims. To say that a moral theory ignores the separateness of persons must not simply be a roundabout way of saying that it places too great a demand on individual agents to sacrifice their integrity for aggregate well-being, or that it countenances the violation of individual moral rights. These might be—indeed, I think they are—perfectly good reasons to object to utilitarianism. But appeals to the separateness of persons seem to promise more than this. They seem to promise an explanation of why utilitarianism has the normative defects—an explanation grounded in its failure to take into account some relevant aspect of our human nature, namely, our separateness. To say that a theory fails to take seriously the separateness of persons is not—or ought not to be—just to repeat, in different language, the stock normative objections to utilitarianism. It should be a way of explaining and thus strengthening those objections by pointing to the metaphysical facts which underlie them.

Properly understood, then, the claim that a theory ought to take seriously the separateness of persons will not be either a purely factual claim or a purely normative claim. It will be both. It will consist of a factual claim to the effect that there is some way in which persons are separate, and a normative claim to the effect that a good moral theory ought to take account of this factual separateness in a certain sort of way. Failing to take the separateness of persons seriously is a normative error about how to respond to a certain fact.

3.

But what sort of normative error is it? Judging by the way in which both Rawls and Nozick state the problem, the main worry seems to be that there is some feature of classical utilitarianism which leads to that theory's sanctioning intuitively unfair trade-offs. But what feature is it? And how is this feature related to the descriptive fact of the separateness of persons?

Perhaps the most notable feature of utilitarianism, in this context, is its consequentialistism. A theory is consequentialist, let us say, if it judges the rightness and wrongness of an act solely by virtue of the consequences that act produces or is likely to produce. Classical utilitarianism is, of course, consequentialist, insofar as it defines right acts as those which maximize the net amount of satisfaction in the world. But it would be easy to generate non-utilitarian consequentialist theories simply by altering utilitarianism's conception of the good. Thus, instead of maximizing satisfaction, we might seek to maximize wealth, or longevity, or hair.

It seems clear, however, that the utilitarian cannot escape concerns about failing to take the separateness of persons seriously merely by altering her standard of value in this way. A theory which held that we ought to maximize wealth, or longevity, or even the non-violation of moral rights,[3] seems just as susceptible to the objection as does classical utilitarianism. For under any such theory, all it would take to entirely justify my loss of some good is that someone else experience a large enough gain. If the appeal to the separateness of persons carries force no matter what the good being

maximized is, perhaps utilitarianism's problem lies not in its theory of the good, but in the fact that it takes the goodness of the resulting state of affairs to be the sole criterion for judging the rightness or wrongness of the action. Perhaps, in other words, it is the consequentialism of such theories to which those who appeal to the separateness of persons object.

This interpretation of the separateness of persons is not without its intuitive plausibility, nor is it entirely without precedent. David Brink, for instance, has written that it is the "consequentialist or teleological" structure of utilitarianism which appeals to the separateness of persons is meant to target, and thus concludes that all theories which share that structure are equally vulnerable to the appeal.[4] Appeals to the separateness of persons are thus not limited in scope to classical utilitarian theories, nor even to utilitarian theories more broadly construed. Any theory which judges the rightness and wrongness of acts solely on the basis of the goodness or badness of the consequences they produce is equally vulnerable.

As plausible as this interpretation may seem, I believe it is subject to a decisive objection. At the beginning of this paper, I framed the worry lurking behind the separateness of persons as one pertaining to trade-offs. In so doing, I was following the exposition which the problem has been given in Rawls, Nozick, and most other proponents of the objection. To reiterate, the concern seems to be that utilitarianism allows certain individuals to be sacrificed without limit for the benefit of some aggregate good. But what leads to this result is not merely the theory's defining the right in terms of the production of good. Rather, it is the fact that theories like utilitarianism ask us to maximize the good *as a whole,* irrespective of the effect of that maximization on particular agents. But there is nothing in consequentialism, as we have defined it, which commits it to this. A theory is consequentialist if it determines the rightness of acts based on the consequences of those acts. Such a consequentialist theory need not judge the goodness of consequences on the degree to which some agent-neutral good has been maximized. For instance, an egalitarian theory which holds those acts are right which tend to produce the most nearly-equal distributions is consequentialist, but does not appear to run into the same problem with trade-offs that classical utilitarianism does. Such a theory would sanction taking from one individual to give to another, but does not, in doing so, seem to run afoul of the separateness of persons in the way that classical utilitarianism does. Classical utilitarianism supports such trade-offs if and to the extent that they maximize the agent-neutral good of happiness, regardless of how particular individuals fare under the resulting distribution. Egalitarians, on the other hand, support

Matt Zwolinski

trade-offs precisely because of considerations about how particular persons are faring—the entire point is to ensure that different people fare as close as possible to equally well. While classical utilitarianism can thus plausibly be said to ignore the separateness of persons, egalitarianism seems to be premised on the importance of the separateness of persons.

Thus, if concern with the legitimacy of certain trade-offs is fundamental to the separateness of persons, not all consequentialist theories seem equally vulnerable to that criticism. Specifically, consequentialist theories which are distribution-sensitive will seem less susceptible. To say that a consequentialist theory is distribution-sensitive is simply to say that it bases its ranking of states of affairs at least in part on the way in which whatever value the theory seeks to promote is distributed. Such theories might assign "bonus points" to those distributions which most closely resemble equality, or might specify that the resulting gains and losses of the least well off members of society are to be given more weight in the evaluative calculus than those who are the better off. Assuming that a sufficiently large multiplier is applied to those distributions that resemble intuitively fair distributions, a distribution-sensitive consequentialist theory thus might not be vulnerable to the charge that it allows for unfair trade-offs. In output, if not in terms of rationale, it seems possible that such theories could be made to be equivalent to those favored by the critics of consequentialism.

4.

Perhaps, then, it is not the consequentialism, but the distribution-insensitivity of classical utilitarianism which renders it vulnerable to appeals to the separateness of persons. It is fairly easy to see how this argument might work. Let us return to the analogy of trade-offs. On many theories of rationality, an individual ought to be neutral among the various ways in which the good can be distributed across his life, temporally speaking. The point, for such theories, is to have as good a life as possible. Insofar as different distributions can affect how well one's life goes as a whole, they are derivatively relevant. But temporal distribution as such has no fundamental importance. Rational agents should be *temporally-neutral.* The utilitarian, however, seems to move from the rationality of temporal-neutrality to the rationality or desirability of *person-neutrality.*[5] Just as distribution across separate time-slices does not matter within a life, so, the utilitarian holds, distribution among separate persons should not matter across lives. The social decision procedure is, in this sense, no different from the individual

decision procedure. In each case, the component parts of the whole among which the good can be distributed are viewed as fundamentally irrelevant. Each decision procedure concerns itself with the whole—the whole life or the whole society—rather than the parts of which that whole is composed. Insofar as person-neutral theories do not ask or care how the good is distributed among distinct individuals, they can plausibly be said to fail to take seriously the separateness of persons.

If person-neutrality is the feature that makes a theory vulnerable to the charge of failing to take separateness seriously, then the set of vulnerable theories will, as we suspected earlier, be broader than classical utilitarianism alone. However, as we have seen, it will also be smaller than the set of consequentialist theories, for not all consequentialist theories are person-neutral. Some are distribution-sensitive.

Distribution-sensitivity can, of course, be built in to a theory in a number of different ways. At one extreme, a theory might forbid trade-offs entirely, holding that the only morally justified acts are those which yield a pareto-superior distribution of the good.[6] Most theories, however, would likely allow a fairly wide range of trade-offs. Some, such as those which assign a heavier weighting to equality-approximating distributions, seem likely to do so in a way which rules out many of the more objectionable trade-offs countenanced by classical utilitarianism. Others, however, might still allow for trade-offs which strike us as intuitively objectionable. Such theories might, for instance, assign a heavier weighting to the well-being of the least well-off, but the weight might not be high enough to rule out public policies which cause a small underclass to suffer in relative misery for the greater benefit of the upper and middle classes. Such theories might thus still allow for intuitively unfair trade-offs. Are theories which allow for such results just as bad as classical utilitarianism at respecting the separateness of persons?

At least on the current interpretation, the answer seems to be, "No." After all, the interpretation we are considering holds what it *means* to ignore the separateness of persons is to be a person-neutral theory. And the sort of theory we are now considering is not person neutral. It is distribution sensitive, insofar as the way in which the good is distributed is of non-derivative importance to the way in which the theory evaluates states of affairs. If appeals to the separateness of persons are targeted simply at those theories which ignore matters of distribution entirely, then *any* distribution-sensitive theory will be immune to the criticism, no matter how small the weighting it gives to the well-being of the least well off, or to near-equal

distributions, and so forth.

If we understand the objection as applying only to person-neutral theories, then, it seems that we will be unable to apply it to a wide range of theories which seem just as intuitively objectionable as utilitarianism. The simple solution, then, seems to be that we should recharacterize the objection as applying not only to those theories which ignore matters of distribution, but those which fail to give such matters sufficient weight in their evaluative calculus. It is not enough simply to assign some weight to matters of distribution in one's theory. If the weighting one gives is not enough to rule out the sort of objectionable trade-offs countenanced by classical utilitarianism, then the theory can still be said to ignore the separateness of persons.

This seems to be the most natural and intuitive interpretation of appeals to the separateness of persons that we have considered thus far. Nevertheless, some very serious worries still remain. Things begin to get tricky when we are faced with deciding how much weight to give to matters of distribution. Should the interests of the least well off be given lexical priority over the needs of all others in society? Or should they only be weighted, say, twice as heavily as the interests of the median person? There might be good reason for favoring one answer to this question over another—reasons of public policy, for instance, or of other moral considerations such as liberty or justice. But can we really say that one answer is better than another at taking the separateness of persons seriously? A theory which considers the suffering of the least well off to be insufficient to justify redistributing resources to her from a better off individual might be said to be failing to take her suffering seriously enough. But is it really failing to take her separateness seriously?

Things are even more complicated than this, for we must decide not merely how much weight to give to matters of distribution, but which distributive facts we wish to assign weight to in the first place. As we saw earlier, there is an extremely large number of ways in which a theory can be distribution sensitive. Which of these can be said to properly recognize the separateness of persons? One which assigns "bonus points" to distributions which approximate equality? One which maximizes the well-being of the least well off member of society? Even egoist theories are distribution sensitive, but here the problem seems to be that the theory gives *too much* weight to the separateness of persons. Which distributive facts are important to recognize if we want our theory to take seriously the separateness of persons?

The problem here is not deciding which criterion for judging distributions is intuitively the most just or the most fair. The problem is to provide some grounding for our intuitions of fairness and justice by recourse to the separateness of persons. Some distribution-sensitive theories yield intuitively fair distributions, and others do not. But if all we have to rely on are our intuitions of fairness and justice, then it looks as though appeals to the separateness of persons are doing no real work. For such appeals to have any real bite, we should be able to point out that those distributions which strike us as intuitively more fair or just are also those distributions which better respect the separateness of persons in some straightforward sense. But so far, we have encountered no argument as to why egalitarian theories better respect separateness than inegalitarian theories. Why should we believe that theories which require that the well-off be sacrificed in order to serve the poor better respect separateness than those theories which require the poor to sacrifice for the well-off? Almost any plausible distribution-sensitive theory will sometimes require that certain individuals forgo some benefit, or undergo some harm, for the benefit of others often without any hope of compensation. If whatever normative conclusion we are supposed to draw from the separateness of persons is not so strict as to rule out all such trade-offs (as surely it could not plausibly be), then what guidance can it possibly give us as to which trade-offs are permissible and which are not?

5.

Our conclusions seem to bode ill for the separateness of persons. The main result of our inquiry thus far is that we have been unable to isolate any feature of classical utilitarianism which could plausibly be held responsible for both its sanctioning of intuitively unfair trade-offs and its failure to give proper recognition to the ways in which persons are factually separate. The distribution-insensitivity of the theory seems to come close to the mark, but not just any sort of distribution-sensitivity will solve the problem, and it is far from clear how appealing to the mere fact of separateness alone can do anything to resolve the debate between competing theories of how distribution-sensitivity should be instantiated. Furthermore, even if we were to have identified such a feature—even if we could pick out that component of utilitarian theories which was responsible for its sanctioning of intuitively unfair trade-offs—this would still be a long way from providing an *argument* against the theory. After all, the argument is not that the utilitarian has

ignored some fact of separateness in a straightforward way. It is that she has failed to give that fact the proper normative weight in her theory. But surely, the utilitarian thinks she has given precisely the proper normative weight to the fact of separateness—zero. She may, of course, be wrong about this. But demonstrating that she is wrong requires argument, and the sorts of appeals philosophers have made to the separateness of persons do not seem to provide such argument. They thus appear to be not so much an explanation of utilitarianism's flaw, as an expression of the idea that there is some flaw which, presumably, further argument will reveal.

A proper argument based on the separateness of persons would, as I stated earlier, have to do two things. It would have to point out some feature of classical utilitarianism that is responsible for its objectionable entailments on issues of distributive justice, and it would have to show how this feature either commits utilitarianism to, or is the result of, failing to give the proper consideration to some way in which persons are metaphysically separate. I have attempted to show that none of the most plausible ways of fleshing this theory out are successful. But I do not think we should conclude from this that appeals to the separateness of persons are entirely without philosophical merit. I believe that a good, anti-utilitarian argument can be made based on considerations of the separateness of persons, but it would require an approach altogether different from the ones we have so far examined.

I think it is probably a mistake to view intuitively objectionable trade-offs as the prime example of utilitarianism's disrespect for the separateness of persons. That it sanctions unfair trade-offs is a problem for utilitarianism, but I think utilitarianism's attitude toward trade-offs is merely the result of a deeper problem with utilitarianism that is better characterized as a lack of respect for separateness. That problem has been described by Bernard Williams as utilitarianism's inability to respect an agent's integrity.[7] Utilitarianism, in demanding agents to devote all their energies to maximizing agent-neutral good, shows insufficient respect for an agent's commitments to the projects and goals that give her life meaning. There is, I think, much work to be done in developing this objection and showing how this feature of utilitarianism is based on its failure to respond properly to some aspect of our metaphysical separateness as persons.[8] Still, I think that this way of understanding appeals to the separateness of persons has much more to be said in its favor as an explanation of utilitarianism's shortcomings than do those which attempt to locate them in its person-neutral or consequentialist structure.

NOTES

[1] David Gauthier, *Practical Reasoning: The Structure and Foundations of Prudential and Moral Arguments and their Exemplification in Discourse* (Oxford: Oxford UP, 1962) 128.

[2] Robert Nozick, *Anarchy, State, and Utopia* (New York: Basic Books, 1974) 32-33.

[3] This is akin to Nozick's "utilitarianism of rights." See Nozick (1974). While Nozick rejects this theory, he does not explicitly do so on the grounds that it ignores the separateness of persons. Nevertheless, the theory seems subject to the same basic type of objections which Nozick employed against classical utilitarians as an elaboration of his appeal to the separateness of persons. An individual whose rights are violated does not gain any compensation from the fact that the aggregate level of rights-violation in society is lower as a result, for she has only one life to live, etc.

[4] David Brink, "The Separateness of Persons, Distributive Norms, and Moral Theory," in Frey and Morris, eds., *Value, Welfare and Morality*. (Cambridge: Cambridge UP, 1993) 253.

[5] See Brink 253.

[6] Distribution *A* is pareto-superior to distribution *B* just in case one member of the group is better off under *A* than *B*, and no member of the group is better off under *B* than *A*.

[7] Bernard Williams, "Persons, Character and Morality," reprinted in his *Moral Luck: Philosophical Papers 1973-1980* (Cambridge: Cambridge UP, 1981).

[8] I attempt to engage in some of this work myself in *Morality's Authority*, chs 2-3 (manuscript in progress).