Trim: 228mm × 152mm Top: 12.653mm Gutter: 21.089mm
CUUK2041-10 CUUK2041/Russell ISBN: 978 1 107 00116 9 August 25, 2012 12:18

MATT ZWOLINSKI AND DAVID SCHMIDTZ

10 Environmental virtue ethics
What it is and what it needs to be

ACT AND CHARACTER, PRINCIPLE, AND RULE

If we ask what makes an action right, one plausible answer is that the right action is the one that does as much good as possible. Roughly speaking, this is the theory known as consequentialism. The theory is most often associated with John Stuart Mill, and it is one of the simplest theories we have. An alternative theory: What makes an act right is not whether it *promotes* what is good so much as whether it *respects* what is good. Associated most often with Immanuel Kant, this theory is known as *deontology* and says, more specifically, that the only thing that is an unqualified good in itself is the good will of an autonomous person and therefore, an action is right if, but only if, it expresses respect for all persons as ends in themselves and therefore treats no person merely as a means to further ends.

Yet another alternative, *virtue ethics*, is so different it might be best to see it not as an alternative answer to the same question but as responding to a different question altogether. Often associated with Aristotle, but with roots in various traditions as discussed in this volume, virtue ethics tells us that what is right is to be a certain kind of person, a person of virtue: courageous, modest, honest, even-handed, industrious, wise. A virtuous person will, of course, express his or her virtue through action. But, for virtue ethics, the specification of rules of right action is largely a secondary matter – one that in many ways presupposes the kind of practical wisdom possessed by the person of virtue.

This chapter incorporates material from Zwolinski and Schmidtz 2005, and also from Schmidtz 2001a.

Trim: 228mm × 152mm Top: 12.653mm Gutter: 21.089mm
CUUK2041-10 CUUK2041/Russell ISBN: 978 1 107 00116 9 August 25, 2012 12:18

We wish we could settle which of these theories is right, then specify what that correct theory tells you to do. For better or worse, though, moral life is more complicated than that. The three theories just described are the main theories we discuss in introductory classes in moral philosophy, but most philosophy professors believe that none of them represents the single best way of capturing all that is true about morality. Morality is complex. It calls for creativity and judgment in the same way that chess does. You may come to the game of chess hoping to learn a simple algorithm that picks out the winning play no matter what the situation. For human players, though, there is no algorithm. There is no substitute for creativity and good judgment, for the ability to think ahead and expect the unexpected. Even something as simple as a game of chess is full of surprises, yet the complexity involved in playing chess is nothing compared to the complexity involved in being moral.

One could think, and we do think, that this way of understanding the challenge of being moral is most at home in a virtue ethical approach. Our students seem often to come to their first ethics course hoping to be given a list of rules or a code of professional conduct. Experience tells us, though, that when moral philosophers try to do applied ethics, there turns out to be something artificial and unhelpful about trying to interpret morality as a set of rules. Rules function in our reasoning as trump cards. If we have a rule, and can believe with complete confidence that the rule ought to be followed, and if we ascertain that a certain course of action is forbidden by the rule, that settles it. The rule trumps all further reasoning, so no further reasoning is necessary.

How comforting it would be to have such rules. And of course, sometimes the situation actually is rule governed. Not always, though. Often, there are reasons favoring an action, and reasons against, and neither trumps the other.

It may still be possible, however, to decide in a principled way. *Principles* are not like rules. Where rules function in our reasoning like trump cards, principles function like weights. If the applicable moral rule forbids X, then X is ruled out, so to speak. In contrast, principles can weigh against X without categorically ruling out X.

Consider an analogy. A home builder might say, in describing his or her philosophy about how to build houses, "You have to minimize ductwork." Question: Is that a rule or a principle? The answer is that,

Trim: 228mm × 152mm Top: 12.653mm Gutter: 21.089mm
CUUK2041-10 CUUK2041/Russell ISBN: 978 1 107 00116 9 August 25, 2012 12:18

interpreted as a rule, it would be silly. As a rule, it would say, no matter what weighs in favor of more extensive ductwork, minimize ductwork, period. In other words, zero ductwork!

In fact, though, "minimize ductwork" is a good principle rather than a bad rule. As a principle, it tells home builders to be mindful of energy wasted and living space consumed when heated or cooled air is piped to remote parts of the house. Other things equal, get the air to where it has to go by the shortest available route. This principle will seldom outweigh the principle that the ceiling should be a minimum of seven feet above the floor. That is to say, it is not a trump, but it does have weight. A good builder designs houses with that principle in mind, but does not treat the principle as if it were a rule.

When students sign up for introductory courses in ethics, some of the most conscientious come in hoping to learn the moral rules. It is a shock when we say we have been teaching ethics for thirty years, but for the most part, we don't know the moral rules, and we suspect there are too few to give comprehensive guidance regarding how we ought to live.

When making real-world practical decisions, the considerations we bring to bear are more often principles than rules. So why, when we look to moral philosophy, would we hope to be given rules rather than principles? What is the attraction of rules? The idea of following a rule is comforting because it has the feel of relieving us of moral responsibility. If we follow the rules, it seems to guarantee our innocence. Unlike rules, principles offer no such escape. Rules are things we follow. Principles are things we apply. There is no illusion about principles being the kind of thing we can hide behind. Principles leave us with no doubt as to who is responsible for weighing them, for making choices, and for bearing the consequences.[1]

The upshot, and it is fundamental to understanding what being a moral agent is like in the real world: if you need to figure out what to do, don't look for rules; look for principles. Needless to say, this too is a principle, not a rule. It has exceptions. There are, after all, rules. Rules sometimes do trump all other considerations.

None of this, we believe, is exactly entailed by the virtue ethical approach, but neither is any of this contradicted by that approach.[2] Virtue ethics is about understanding the challenge of being moral as in the first instance a challenge of being a certain kind of agent, a

Trim: 228mm × 152mm Top: 12.653mm Gutter: 21.089mm
CUUK2041-10 CUUK2041/Russell ISBN: 978 1 107 00116 9 August 25, 2012 12:18

certain kind of character. But part of being a moral character is being principled, and so too being a moral character involves being sensitive to and respectful of morally justified rules, and of principles, when they bear on what is required of a moral agent in the case at hand (Schmidtz 2001a).

REPUGNANT CONCLUSIONS: A FUNDAMENTAL CHALLENGE FOR ENVIRONMENTAL ETHICS

We have outlined a fairly expansive way of understanding virtue ethics, and it is on this ecumenical understanding that Thomas Hill, one of the trailblazing Kantians of our time, transformed the field of environmental ethics with an essay that inspired the emergence of what we now call environmental virtue ethics.[3]

In a classic article, Hill showed how traditional (rights-based) deontological and (utilitarian) teleological theories can fail to track our moral intuitions regarding environmental issues (Hill 2005). When Hill's neighbor cut down a beautiful old avocado tree and covered his yard with asphalt, Hill was indignant, but paused to wonder whether he had any theoretical justification for his indignation. The problem was not that trees have rights, or that Hill's neighbor had no right to cut down the tree. Hill's neighbor was depriving others of the enjoyment of the tree, but although this was consistent with Hill's indignation, it did not really account for it.

Ultimately, Hill concluded, the core question was not what was wrong with the act, but what was wrong with the person. "What sort of person," Hill asked, "would do a thing like that?" Hill's answer was that an admirable person would not do it. A person would have to be insensitive, and lacking in humility. Interestingly, the humility about which Hill was talking was recognizably an ecological humility. It is what Aldo Leopold meant when he spoke of our need to acknowledge and ultimately cherish our proper status as citizens (not conquerors) of the biotic community (Leopold 1966).

Like Tom Hill, many of us who are disturbed by the callous treatment of our natural environment would feel uncomfortable arguing that trees, brooks, or sand dollars have a *right* to be left alone. Also like Hill, though, neither is our discomfort grounded in a calculation of aggregate utility. The problem is that both deontological and

Trim: 228mm × 152mm Top: 12.653mm Gutter: 21.089mm
CUUK2041-10 CUUK2041/Russell ISBN: 978 1 107 00116 9 August 25, 2012 12:18

consequentialist theories locate the source of our discomfort in the wrong sorts of considerations. As a result, they tend to give us false guidance in a variety of situations.

One particularly notorious example of this is Derek Parfit's Repugnant Conclusion (Parfit 1984). Parfit conceived his Repugnant Conclusion as a problem for one version of a utilitarian value theory, but in the sections that follow, we generalize Parfit's argument. We suggest, first, that all utilitarian theories of value face analogous problems, and second, that we cannot solve the problem (although we might be able to obscure it) by adopting a different kind of *act-centered* theory. Ultimately, Repugnant Conclusions suggest (we do not claim they entail) problems for the whole idea that moral theorizing should culminate in a simple formula for right action. We need a different sort of theory. We need not merely a better formula, but a better objective, such as is hinted at in Tom Hill's environmental virtue ethics.

The Repugnant Conclusion

Standard versions of the principle of utility say something like this: an act is right if and only if it maximizes happiness.

What does it mean to maximize happiness? First, the principle is referring to an aggregate: that is, the sum of everyone's happiness.[4] Second, the principle on its face is quantitative, referring to the kind of thing that can be maximized. Third, the quantity is most naturally thought of as a total sum, as opposed to an average. As Parfit notes, however, there is a problem with this seemingly innocuous third point.

Suppose we are deciding whether to have one or two children, and have no reason to doubt that the two children would each be about as happy as the one. If the two children would each be as happy as the one, then we conclude there is more total utility, indeed about twice as much, in having two children than in having one.

On its face, we seem to have utilitarian grounds for having two children rather than one. However, as Parfit points out, when we endorse the principle that the right act maximizes *total* happiness, we commit ourselves to the Repugnant Conclusion that

Trim: 228mm × 152mm Top: 12.653mm Gutter: 21.089mm
CUUK2041-10 CUUK2041/Russell ISBN: 978 1 107 00116 9 August 25, 2012 12:18

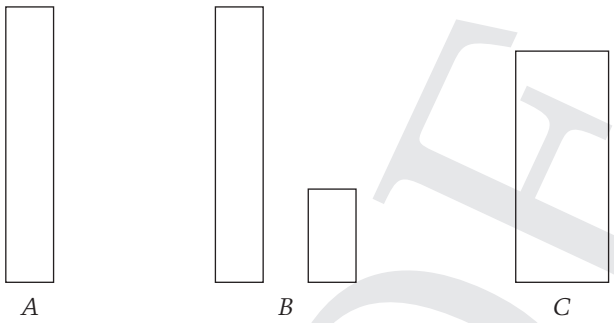

Figure 10.1 Population size and average quality of life. Population *B* differs from *A* only by the addition of a group of persons whose lives are less fortunate, but still worth living. *C* is the result of equalizing the well-being of the two groups in *B* by imposing a small cost on the better-off group for a large gain to the worse-off group. Width of the blocks = population size; height, average quality of life.

For any population of at least ten billion people, all with a very high quality of life, there must be some much larger imaginable population whose existence, if other things are equal, would be better, even though its members have lives that are barely worth living. (Parfit 1984, p. 388)

Here is how Parfit reaches that conclusion. Figure 10.1 illustrates three possible populations. The width of the blocks represents the number of people living; the height represents their average quality of life. In population *A*, people's lives are, on average, well worth living. Population *B* differs from population *A* only in that, in addition to the population of *A*, it contains an added group. These people's lives are worth living, too, though less so than the lives of the persons in population *A*. Let us stipulate that their existence does not affect the persons in population *A* at all – members of population *A* do not even know of this new group's existence. The question, then, is this: do we make a situation *worse* by moving from *A* to *B* – that is, by the mere addition of persons whose lives are worth living?

It's hard to imagine condemning such an addition. By hypothesis, no one's rights are violated, total utility is higher, and each new person is happy to be there.[5] At the very least, *B* seems *no worse* than *A*. How does *B* compare with *C*? In *C*, the better-off group from *B* has been made worse off. But their loss is smaller than the

Trim: 228mm × 152mm Top: 12.653mm Gutter: 21.089mm
CUUK2041-10 CUUK2041/Russell ISBN: 978 1 107 00116 9 August 25, 2012 12:18

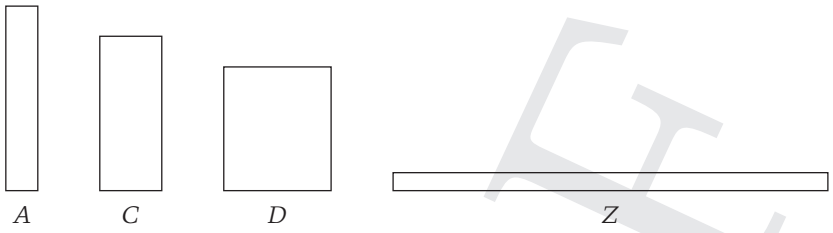

Figure 10.2 Quality of life and increasing population size. Repeating the sequences described in Figure 10.1 leads to a state of affairs *Z* that much worse than the original (*A*). Width of the blocks = population size; height, average quality of life.

gain achieved by the worse-off group, so total utility has risen. The resulting distribution is superior on egalitarian grounds as well. On various grounds, then, *C* is at least as good as *B*.[6]

The problem, of course, arises in the repetition of this sequence, as illustrated in Figure 10.2. For if the move from *A* to *C* is justified, why not the move from *C* to *D*, and so on, all the way to *Z*? In other words, if a state of affairs is made better (or at least, not worse) than another by doubling the population while decreasing average utility by less than 50 percent, why not continue until we are left with an *enormous* number of people whose lives are, on average, only *barely* worth living?[7]

This is the Repugnant Conclusion. To people seeing this puzzle for the first time, the conclusion that *Z* is morally desirable, or even that it is merely no worse than *A*, sounds absurd. How could a society be superior to another simply in virtue of being so much larger, when average members are so much worse off that *one more* traffic jam, stubbed toe, or malfunctioning toilet will drive them to conclude that life is no longer worth living?

Generalizing: a problem for utilitarianism

Here is an easy response. Total utilitarianism works well enough in a world of fixed population, but in a world where population size is one of our choice variables, we need a more sophisticated metric. In this more complex world, the intuitive attraction of utilitarianism is better captured by *average* utilitarianism: the theory that an act is

Trim: 228mm × 152mm Top: 12.653mm Gutter: 21.089mm
CUUK2041-10 CUUK2041/Russell ISBN: 978 1 107 00116 9 August 25, 2012 12:18

right if and only if it maximizes average utility. As Parfit was aware, though, average utilitarianism has its own version of the problem.

The Other Repugnant Conclusion: For any population of people with a quality of life more or less like ours, say, we can imagine a smaller population (in the limit, a single, godlike Utility Monster) that is on average so much happier that it would be better if our population were replaced by that smaller one.

Some environmentalists will not find this repugnant, but this is because they have other reasons to oppose overpopulation. To them, there are things in the world that matter more, maybe a lot more, than the happiness of persons. In other words, they are not utilitarians. To utilitarians, though, the *Other* Repugnant Conclusion is almost as big a problem as the original.

To summarize, the problem in its general form is that we have two kinds of mechanical measures of aggregate happiness: total or average. In practice it may be massively difficult, indeed impossible, truly to arrive at any such measure. But the problem suggested by our Repugnant Conclusions goes deeper: that is, we would not be able to trust a *number* even if it were easy to obtain, and even if its accuracy were indubitable. The fact would remain that neither version of a utilitarian number – that is, neither total nor average happiness – reliably tracks the intuitions that led us to find utilitarianism plausible in the first place.[8]

Generalizing further: a problem for all act-centered theories

The problem is more than a problem for utilitarianism. It is a problem for all *act-centered* theories: that is, moral theories that treat specifying action-guiding rules as their primary task. From an act-centered perspective, it is hard to explain *why* the Repugnant Conclusion is repugnant. We have already seen that utilitarian theories, which base their evaluation of actions on their tendency to produce desirable states of affairs, are left with little means of stopping once they accept the desirability of the initial moves from *A* to *B*, and *B* to *C*.

Deontological theories might appear to be in better shape. After all, they do not base their evaluation of an act's morality on its tendency to produce a certain state of affairs, so they can reject the move

Trim: 228mm × 152mm Top: 12.653mm Gutter: 21.089mm
CUUK2041-10 CUUK2041/Russell ISBN: 978 1 107 00116 9 August 25, 2012 12:18

from *A* to *B* even if they judge *B* to be a better state of affairs. But it is an awkward response for a theory to bite the bullet and say (as the hardest forms of deontology do) that consequences should play *no* role in our evaluations of states of affairs. But that consequences should play some role is all that is needed to get Repugnant Conclusions off the ground. For if *C* is superior to *A* on *some* morally relevant grounds, then the burden of proof is on the deontologist to show that there are countervailing moral considerations that override the moral case for moving from *A* to *C*. And this seems unlikely. Who is being treated as a mere means in the course of moving from *A* to *C*? Where in the chain is the step that cannot be universalized? For that matter, what if anything is wrong with *Z* from a deontological perspective? Whose rights are being violated? More profoundly, whose rights are violated by the bare admission that *Z* is a better state of affairs? A theory might try to save face by insisting that the move from *A* to *Z* is for some reason impermissible, but quibbling about the propriety of the move would be too little too late. If a theory admits that *Z* is *better*, it has already embarrassed itself, regardless of whether it prohibits *moving* to *Z*.[9]

Deontological theories are more apt for considering how to treat currently existing people than for considering whether it would be good for an additional population to come into existence. Which is to say, deontological theories are less apt for a world where moral problems are increasingly taking on ecological dimensions.

Somehow, the initial utilitarian rationale for the move from *A* to *C* is spurious. In some way, a total utilitarian has the wrong idea, not about whether consequences are relevant, but about the way in which consequences are relevant.[10] At best, an average utilitarian does only a little better. We may hope for some escape from the problem once we factor in environmental considerations, but it is built into the problem that life in the larger population is, after all, worth living. So whatever loss of environmental amenities people face, average people nevertheless are eking out lives worth living, even if only barely.

An anthropocentric deontology likewise would seem not to solve the problem, as there is no particular reason to suppose members of this larger population are failing to treat each other as ends in themselves. A deontology expanded to embrace animal rights merely treats one symptom of a larger problem, because the

Trim: 228mm × 152mm Top: 12.653mm Gutter: 21.089mm
CUUK2041-10 CUUK2041/Russell ISBN: 978 1 107 00116 9 August 25, 2012 12:18

problem goes beyond our treatment of other sentient creatures to the wanton destruction of trees and such. The larger concern is not animal rights; it is more fundamentally an ecological concern, perhaps with aesthetic overtones.[11]

IS THERE ANOTHER WAY OF DOING MORAL THEORY?

There would seem to be a more direct way of going to the heart of the problem, but it would involve giving up on standard act-centered moral theories. The heart of the task is at least in part to define ideals of human excellence – to define a conception of the good life for a person that makes sense as an ideal, with no presumption that the ideal must be adopted as a goal to be promoted. We have other ways of responding to ideals; simply respecting them comes to mind, an attitude perhaps more in keeping with intuitions that motivate deontology. Part of an ideal of human excellence could, in turn, incorporate forms of environmental sensitivity. Again, though, there is no assumption that environmental sensitivity must translate into an activist agenda – simply appreciating the beauty of nature is among the more admirable ways of being sensitive to it.[12]

Thomas Hill suggests that our discomfort with environmental exploitation is not wholly a product of our belief that environmental goods are being put to inefficient use, nor that those who exploit such goods are violating any rights in doing so. It is a mistake, he thinks, to suppose that "all moral questions are exclusively concerned with whether *acts* are right or wrong, and that this, in turn, is determined entirely by how the acts impinge on the rights and interests of those directly affected" (Hill 2005, p. 48). Instead, Hill suggests, we ought to ask "What sort of *person* would destroy the natural environment?" Approaching the issue from this perspective allows us to see that "even if there is no convincing way to show that the destructive acts are wrong (independently of human and animal use and enjoyment), we may find that the willingness to indulge in them reflects the absence of human traits that we admire and regard as morally important" (Hill 2005, p. 50). People who carve their initials in 100-year-old Saguaros might not be violating any rights, and the satisfaction they get might well outweigh the suffering caused to other sentient beings, but the fact remains that there is some defect in such people's characters. In Hill's language, they lack a kind of

Trim: 228mm × 152mm Top: 12.653mm Gutter: 21.089mm
CUUK2041-10 CUUK2041/Russell ISBN: 978 1 107 00116 9 August 25, 2012 12:18

humility – an ability to appreciate their place in the natural order or to "see things as important apart from themselves and the limited groups they associate with" (Hill 2005, p. 51).

Hill's approach offers a natural, straight-to-the-point way of thinking about what Richard Routley called the "Last Man Argument." Routley's thought experiment presents you with a situation something like this: You are the last human being. You shall soon die. When you are gone, the only life remaining will be plants, microbes, invertebrates. For some reason, the following thought runs through your head: "Before I die, it sure would be nice to destroy the last remaining Redwood. Just for fun." What, if anything, would be *wrong* with destroying that Redwood? Destroying it won't hurt anyone, so what's the problem? The problem is, what kind of person would destroy that last Redwood? What kind of person would enjoy such wanton destruction of such a beautiful, majestic, living thing? Hill's question seems like exactly the right question.[13]

Indeed, for a number of philosophers Hill's question promises to transform the way we think about environmental ethics. The relatively recent field of environmental virtue ethics has developed in order to study the norms of character that ought to govern human interaction with the environment, and scholars such as Philip Cafaro, Jason Kawall, and Ronald Sandler are developing this field in a way that is yielding fruitful theoretical and practical results (Cafaro 2004; Kawall 2003; Sandler 2007). These scholars seek to characterize attitudes and dispositions that are constitutive of environmental virtue, and to explain the proper role of an ethic of character within a broader environmental ethic. And practically, scholars have focused both on the implications of and prerequisites of particular virtues such as benevolence or temperance, and also on the implications of a virtue ethics approach for specific practical problems such as consumerism and genetically modified crops.

Hill's approach also offers a way of dealing with Repugnant Conclusions. For even if we cannot provide a definitive account of the wrongness of preferring a society *Z*, of the sort described in the Repugnant Conclusion, there remains *something* wrong with being the kind of person who would prefer *Z*. That something could be hard to articulate, but no less real for that. In any case, Hill does provide some pertinent articulation. The sort of person who would prefer *Z* is the sort of person who does not possess the humility that

Trim: 228mm × 152mm Top: 12.653mm Gutter: 21.089mm
CUUK2041-10 CUUK2041/Russell ISBN: 978 1 107 00116 9 August 25, 2012 12:18

would lead a more virtuous person to see value in human society playing an appropriately limited role in the biotic community, for nonanthropocentric as well as anthropocentric reasons.

It is difficult to make accurate judgments of character without setting the context in some detail. And it is difficult to imagine a context in which a person's expressed preference for *Z* is worthy of being taken seriously enough to merit a moral evaluation. In the next section, we will attempt to describe a situation in which it makes sense to speak of a person "choosing" to move from *A* to *Z*. For now, though, suppose that someone you know were offered the option of snapping their fingers and thereby popping into existence a population either like *A* or like *Z* in some far-off and (otherwise) causally isolated universe. Presumably, since no other values are at stake, a disposition to choose population *Z* in such circumstances is simply part of what it *means* to believe that *Z* is a superior state of affairs to *A*.

Given this fact, then, what are we to say about the character of a person who prefers *Z* to *A*? Bear in mind that *Z* is, in the end, a fairly miserable place. People's lives are, it is true, still worth living. But only barely so. Think of how much misery a person can endure while still believing that life is worth living. Now think of a whole world – a very *crowded* world – filled with such people. What kind of person would bring *that* kind of universe into existence when they could just as easily have produced a universe with a smaller number of very happy people? What would be the point? Intuitively, there is *no* point, contra total utilitarianism.

The most natural explanation for such a disposition seems to be a sort of obsession. It's normal to think that the happiness of particular other people is important. It's normal to generalize from this and think that happiness itself is important. It's *maniacal* to think that this abstraction translates into a reason to prefer, over *A*, the concrete misery of world *Z*. Like cases of obsession in general, what seems to have gone wrong here is an extreme inability to grasp the larger context. In this respect, it is not unlike a person who originally pursues cleanliness for the sake of health (and health for the sake of a long, enjoyable life), but ends up cleaning compulsively. What started off as a reasonable principle – be as clean as you can – has been transformed into a manifestly unreasonable rule that trumps all other aspects of a worthwhile life.

Trim: 228mm × 152mm Top: 12.653mm Gutter: 21.089mm
CUUK2041-10 CUUK2041/Russell ISBN: 978 1 107 00116 9 August 25, 2012 12:18

The compulsive pursuit of cleanliness thus ultimately undermines the very value that led them to pursue cleanliness to begin with. There is something similarly wrong in being a person who would think of maximizing a happiness number, as if, from an anthropocentric perspective, what matters most were *happiness* rather than persons themselves (or something about persons other than their happiness, such as whether they achieve excellence).

NOTE ON CONTRIBUTING TO OVERPOPULATION

In the real world, of course, populations are never the product of any individual's choice. Individuals and families do not choose populations; they choose whether or not to have children. Populations emerge only as a (typically unforeseen and unintended) consequence of the combination of many such decisions. The fact that we can raise questions about the character of one who holds a sincere preference for population *Z* might thus seem to have few implications for deciding what to say about people and policies which *actually move us* toward *Z*. Do character-centered ethics have anything to say about these more practical population-related questions?

Insofar as they are not the product of any single individual's choice, undesirable populations can be regarded as an externality along the lines of air pollution. No individual family's decision to limit the number of children they have significantly affects the size or well-being of the overall population, just as no commuter significantly affects the amount of smog in the air by choosing to ride their bicycle to work one day. In either case, however, the aggregate result of many individuals failing to make such decisions is tragic.

Suppose for argument's sake that adding large numbers to the existing population will contribute to an unsustainable (or undesirable) population growth. In that case, raising a large family would amount to a failure to contribute to a sort of public good.[14] If that were the case, achieving a desirable population would require a general policy of restraint on the part of most families, and having a large family in that situation would be to free-ride on the restraint of others. The irresponsibility here needn't be calculated – many families simply won't know (and won't bother to find out) the consequences of having more children than they can afford. The point, however, is that where act-centered theories fall short of explaining the

Trim: 228mm × 152mm Top: 12.653mm Gutter: 21.089mm
CUUK2041-10 CUUK2041/Russell ISBN: 978 1 107 00116 9 August 25, 2012 12:18

indignation we sometimes feel (or explain it in an ad hoc, roundabout, or otherwise unsatisfying way) when confronted with free-riding that does only minute harm, a character-centered theory provides a rich vocabulary for criticizing those who contribute to Repugnant Conclusions as short-sighted, irresponsible, weak-willed, selfish, and so on.

Does virtue ethics help to provide a specifically environmental perspective on what's wrong with people who contribute to overpopulation? (1) A utilitarian can say an act has bad environmental consequences, which counts for something. (2) A deontologist can say an act that has sufficiently bad environmental consequences will as a result be nonuniversalizable, and may also fail to treat other would-be users of the same environmental amenities as ends in themselves. (3) A virtue theorist can say an actor is a bad person qua member of the biotic community, which is something else. A virtue theorist can acknowledge that the third conclusion is true partly because the first one is, but can go on to say there is more to being a good person than to act in a way that has good consequences. A good person is considerate, and therefore *cares* about consequences. A good person is humble, in the sense of seeing himself as a locus of value in a world where there are many loci of value, and recognizing that it is not only humans who can be worthy of appreciation.

INTUITION AND THEORY

One false ideal for moral theory is the idea that the right theory will be simple in the sense of being able to substitute for the wisdom of experience. Part of the point of the Repugnant Conclusion is that wise persons realize that the intuitions leading them to find utilitarianism plausible are not in fact captured by the simple formulas utilitarians sometimes claim to offer.[15]

It is not as if theory is a radical alternative to intuition. A theory is an attempt to capture our intuitions with a simple formula. How could we expect to do that without losing some of morality's intuitive nuance? Of course theories will have counterexamples! It is in our veins as philosophers to continue to test our theories against the intuitions we intend to be articulating with our theories, and of course those articulations will be an imperfect match.

Trim: 228mm × 152mm Top: 12.653mm Gutter: 21.089mm
CUUK2041-10 CUUK2041/Russell ISBN: 978 1 107 00116 9 August 25, 2012 12:18

Theories try to systematize our intuitions, but that is like trying to launch a ballistic missile in a direction such that its simple trajectory will track the more complex trajectory of a guided missile. Counterexamples take the form of showing where the ballistic missile deviates from the guided missile's more convoluted path. This is not to express skepticism about the whole project of moral theorizing, so much as about the more particular assumption that act-centered moral theorizing is the way to go. Act-centered theories are one way of trying to articulate. There is no reason to assume they are the best way. Nothing like that is guaranteed. What is more or less guaranteed – we see no counterexamples on the horizon! – is that act-centered theories will provide imperfect guidance.

We are not presenting this as a knockdown argument against act-centered theory. Our conclusion is that act-centered theory has a certain kind of value, not that it has no value. Virtue ethics reminds us that providing moral agents with a decision procedure covering all possible situations is not the main purpose of moral theories (if it is even a purpose at all). The people for whom moral theories are intended are people already in the midst of living their lives. They come to philosophy hoping it can help them reflect on their lives. A moral theory is successful if it provides that assistance; unsuccessful if it does not. Cases like the Repugnant Conclusion show us that an act-centered theory is not useful as a universal decision procedure. The proper lesson is not that act-centered theories are useless, though, but rather that we are better off treating act-centered theory as the sort of thing from which wise persons can gain insight that is useful, even if limited.

SUMMARY

There are times, as Tom Hill says, when the question is not what is wrong with the act, but what kind of person would do it. The Repugnant Conclusion seems to show that there are cases where the moral problem, even from an act-utilitarian perspective, is not straightforwardly a problem of how to maximize total utility. Our "Other Repugnant Conclusion" seems to show that average utilitarianism does not solve the problem; therefore, even by its own lights, act utilitarianism, the simplest, most mechanical of all moral decision procedures, is not reliable as a mechanical procedure even in

Trim: 228mm × 152mm Top: 12.653mm Gutter: 21.089mm
CUUK2041-10 CUUK2041/Russell ISBN: 978 1 107 00116 9 August 25, 2012 12:18

principle. Moral decisions require *wisdom*, not mere computational power, and there is no simple recipe for wisdom.

It is mere appearance, and misleading appearance, that act-centered theories are better than agent-centered ones at converting moral decision-making from art into science. Prevailing act-centered theories incorporate theories of value that specify some of the considerations to which a wise moral agent will be sensitive. That is their contribution to moral wisdom. It is a significant contribution, but they have not done better than that, and probably never will.

Human rights matter, as do animal rights, whenever they are at stake. Interests matter, when they are at stake. Treating persons as ends in themselves matters, when persons and their ends are at stake. Perhaps universalizability matters in some independent way, but if it matters in some independent way, it probably matters in virtue of what it says about an agent's character. The idea is that to act in a way that you could will to be a universal law is arguably the essence of acting with integrity. That is, when we do that, we are acting from motives that we would not hesitate to make transparent, for all the world to see. If Kant was right, acting in accordance with what one could will to be a universal law was the essence of good will, which (although 'good will' is a notoriously technical notion in Kantian scholarship) appears to be a state of character. A virtue theory might not agree with Kant that good will is the only thing good in itself, but might readily agree that good will is basic, and that unless one gets one's character in order, the other good things in life become ashes.

Finally, talk of rights, interests, and treating persons as ends seems especially apt when we are talking about how to treat persons, or perhaps other sentient beings, and that is what we were talking about in the previous paragraph. But what if the issue concerns a person's relation to an insentient creature such as a Redwood, or to the biotic community as such? Intuitively, Tom Hill's question is *the* question.[16]

NOTES

1 If morality was a really just (or even primarily) a system of rules, we might expect that very clever teenagers could become moral experts simply by reading about and mastering those rules. The fact that we

Trim: 228mm × 152mm Top: 12.653mm Gutter: 21.089mm
CUUK2041-10 CUUK2041/Russell ISBN: 978 1 107 00116 9 August 25, 2012 12:18

think that moral expertise is more a matter of wisdom and experience than memorization is one reason for thinking this view of morality to be mistaken. See Annas 2004, 2006, §A.3; Hursthouse 1999, pp. 59ff.

2 Indeed, as Dan Russell has noted in discussion, this is one point on which virtue ethicists such as Julia Annas, Elizabeth Anscombe, Rosalind Hursthouse, and Philippa Foot present a more or less united front.

3 See also Philip Cafaro's excellent work on Henry Thoreau as a forerunner of environmental virtue ethics (Cafaro 2004).

4 It must seem obvious we will have trouble contriving a number that reliably represents aggregate happiness, but utilitarians have little choice but to insist it can be done. Or, at least, to insist that the ideal of maximization can serve as a useful guide to moral deliberation even in the absence of reliable numbers. We accept the assumption for argument's sake. But see Schmidtz 1992.

5 The average utility of the entire population has, of course, decreased. But the average utility of the original group remains unchanged, and there are reasons to think this should be the relevant consideration. If we judged the morality of an action based on the average utility of the population which exists *ex post*, rather than *ex ante*, then average utility would seem to condone (secretly) killing off anybody who was less happy than average. Surely this is not the way to make the world a happier place.

6 This is not to deny that there are impermissible ways of moving from *B* to *C*, or even that a state of affairs which would otherwise be desirable can be rendered condemnable by being brought about in an impermissible way. So long as we assume this is *not* the case, *C* looks preferable to *B*.

7 What does it mean to say that a life is only barely worth living? On its face, this description is compatible with a life's being pretty miserable. This, at any rate, is the interpretation we operate with for the purposes of this chapter. On reflection, however, the judgment that a life is worth living might be one to which it is difficult to assign any concrete interpretation in the absence of a particular perspective. A person who is *already living that life* might be more disposed to sincerely believe that their life is worth living than a person who is observing the world from the outside, deciding whether or not it would be worth it to live that life (rather than remain in their current state of oblivion). Perhaps there are psychological forces which impel the person living their life to convince themselves that it is worthwhile. Or perhaps the fact that they are situated in a particular point in time is relevant (costs in the past are sunk, and there's always hope for the future). Are these

Trim: 228mm × 152mm Top: 12.653mm Gutter: 21.089mm
CUUK2041-10 CUUK2041/Russell ISBN: 978 1 107 00116 9 August 25, 2012 12:18

reasons to discount the preference expressed by actual individuals? What's the alternative?

8 Michael Huemer (2008) has likewise argued that the Repugnant Conclusion is inescapable within any utilitarian framework, although he treats this as a reason for accepting the Repugnant Conclusion rather than a reason for leaving utilitarianism behind.

9 In the case of what we call the Other Repugnant Conclusion, the deontological prohibition of the crimes that would be involved in making most of humanity disappear is straightforward. Here too, however, we see a problem with using the language of universalizability or respect for persons to explain what is wrong with the kind of environmentalist who thinks the world would be better if virtually all of humanity were to disappear.

10 One other possibility would be to handle consequentialist issues in contractarian fashion, that is, by saying that the utilities in question are not fungible. It is true that for any given person, more utility is better for that person, but it is not true that anyone's utility can be traded for anyone else's in such a way as to produce a higher aggregate. Such a theory would reject what otherwise seems to be a utilitarian case for moving from *A* to *Z*, but at a cost of refusing to entertain comparisons at all, and thus also refusing to acknowledge that *A* is preferable to *Z* in some agent-neutral way. For discussion of nonaggregative forms of utilitarianism, see Coons, forthcoming.

11 For related arguments that recast the question as one about the right thing to do rather than reasons for preferring, see Narveson 1967. See also Booin-Vail 1996.

12 Matthews (2001) argues that militant environmentalists make the same mistake as everyone else who seeks to conquer the biotic community: failing to embrace a truly ecological ethos of "letting it be."

13 This paragraph borrows from Schmidtz and Willott's "The Last Man and the Search for Objective Value," chap. 2 of their 2001.

14 To be clear, we do not accept the premise, but it nicely illustrates how indignation over minute harms can often be better explained in terms of universalizability than in terms of utility, and even better, in terms of what such behavior says about the characters of people who free-ride at other people's expense.

15 One of the many features that made Mill a great moral philosopher was his refusal to place more weight on the simple formula of utilitarianism than it was meant to bear. Readers of *Utilitarianism* will find nothing like an algorithm at work in Mill's thoughtful and nuanced moral analysis.

Trim: 228mm × 152mm Top: 12.653mm Gutter: 21.089mm
CUUK2041-10 CUUK2041/Russell ISBN: 978 1 107 00116 9 August 25, 2012 12:18

16 Schmidtz (2001b) argues that we can respect nature without being species egalitarians. Indeed, the view that potatoes and chimpanzees have equal moral standing is incompatible with genuine respect for nature. Genuine respect acknowledges what living things have in common, but also acknowledges differences.