How to Know What Should Be So:

Ethical Guidance and Ethical Theories

Jason Zarri

1. *Introduction*

If one is in a moral quandary it is wise to look for ethical guidance if one has the time to do so. Ethical theories are, among other things, intended to be one possible source of ethical guidance. If such guidance is valuable, then in ethics there is an embarrassment of riches: There are multiple, well-accepted, yet mutually inconsistent theories. These include utilitarianism, Kantianism, virtue ethics, contractarianism, libertarianism, natural law theory, some forms of moral particularism, and more. The disquieting thing is that, at present, it seems that we are not at all close to being able to determine which of them, if any, is right. How can you know what you should do when ethicists, those who devote their careers to studying such theories, cannot reach a consensus on which one we should accept? Those who look to ethical theories for ethical guidance are apt to be disappointed. This situation is problematic, for if ethical theorizing is to have relevance to real-world ethical behavior, and not just be a way of examining ethical issues out of a love of arguments or puzzles, it must be possible for us to use ethical theories to inform ourselves of what we should do.

2. *A Possible Solution to the Problem*

It seems that philosophers have usually tried to address the issue of how one should act by advancing arguments for or against these theories (or certain parts of them). I want to approach this issue from a different angle. The question I will address is this: Can you get ethical guidance about what you should do in certain situations *without* knowing, or even having good reasons to believe, that any particular ethical theory is right?

I know of at least one philosopher who thinks you can. In the following passage from his article “Hunger, Duty, and Ecology”, which was the inspiration for the ideas I express in this article, Mylan Engel Jr. rebuts an objection to the obligatoriness of donating to famine relief:

> One of the most common reasons that I have heard philosophers give for rejecting the arguments of Singer and company [for contributing to famine relief] runs roughly as follows:
>
> > Singer's preference utilitarianism is irremediably flawed, as are Kant's ethics, Aieken's theory of human rights, and Rawlsean contractarianism. The literature is peppered with devastating objections to these views. Because all of the aforementioned arguments are predicated on flawed ethical theories, all these arguments are also flawed. Until someone can provide me with clear moral reasons grounded in a true moral theory for sending large portions of my income to famine-relief organizations, I will continue to spend my money on what I please.
>
> Such a self-serving reply is both disingenuous and sophistical. It is disingenuous because, as noted earlier, utilitarianism, Kantian ethics, human rights-based ethics, and contractarianism are among the most widely accepted theories in normative ethics. In other contexts, philosophers typically embrace one of these four theoretical approaches to ethics. It is sophistical because a similar reply can be used to "justify" or rationalize virtually any behavior. Because no moral theory to date is immune to objection, one could, for example, "justify" rape on the grounds that all of the arguments against rape are based on flawed ethical theories.
>
> The speciousness of such a "justification" of rape is obvious. No one who seriously considers the brutality of rape can think that it is somehow justified / permissible *simply because* all current ethical theories are flawed. But such specious reasoning is often used to "justify" allowing millions of innocent children to starve to death each year. [footnote omitted] (*Environmental Ethics*, p. 462).

Engel goes on to justify his conclusions about donating to famine relief by appealing to what he takes to be almost universally shared commonsense beliefs about morality.

My approach will be different. My idea is that if you compare all the viable ethical theories that you know of, and find that all, or at any rate a great majority of them agree about whether an action you're considering is right, wrong, or permissible, then you know that it is at least highly probable that that action really is right, wrong, or permissible. For if all ethical theories agree about the moral status of an action, it can only fail to have that status if they are all false. And if a great majority of ethical theories agree about the moral status of an action, it can only fail to have that status if all of the theories that agree about its status are false, which becomes more and more improbable as the number of the theories that agree increases. Note that I'm *not* arguing that if a great majority of ethical theories agree about the moral status of an action then it automatically follows that it very probably has that status. The argument is rather that *if some ethical theory or other* is true, *then* majority agreement implies that the action very probably has the moral status that the majority of theories agree that it has. The upshot is that my approach should be a good guide as to what you should do as long as some ethical theory or other is true. So by using my approach you can be guided by ethical theories without having to attempt the difficult task of determining which of them is right.

To clarify, my idea is not to put different ethical theories together to get a composite theory, but to help someone figure out what they should do in a fairly specific situation. By comparing different ethical theories you might find that they agree about *what* you should do in a situation, but they might not agree about *why* you should do it. If you try to "combine" the guidance you get across diverse situations the result would probably not be cohesive enough to yield an ethical theory. And if you find that a sufficient number of ethical theories don't agree even about *what* you should do in a given situation, I think the most you can conclude is that you don't *know* what you should do, and in such a case you would not have any determinate guidance. So I think that in some cases my approach will give you guidance and in other cases it won't.

Nevertheless, on my view, you would have a reason to explore as many different candidate ethical theories as you can, even though there is no need to determine which of them is right. This is because the more of them you consider, the more certain you can be that you have a representative sample of all the possible viable ethical theories, and the more representative the sample is, the more certain you can be that an action really has the moral status you think it does given that the majority of ethical theories agree that it has it.

It is important to note that ethical *theories* properly so called need not be the only kind of ethical view that one might have to take into account on my approach. Timothy Chappell has introduced the different yet related notion of an *ethical outlook*, which he characterizes as follows:

> Anybody who is going to live a genuinely worthwhile and a fully human life will have to live out a set of views and commitments about the central questions concerning value: what is worth living for and what is worth dying for, what is really admirable and what is really contemptible, what we must do at all costs and what we must not do no matter what; and so on. This set of views and commitments need not be very explicit; but it must run deep—must be sincerely and indeed passionately held. And it need not be very systematic; but it must be as considered, rationally defensible, and coherent as possible. Any such set of views about value is what I will call an *ethical outlook*.[1]

The notion of an ethical outlook is, in principle, broader than the notion of an ethical theory, for it can include ethical theories as well as ethical views that are less systematic. In Chappell's opinion, mainstream ethical theories[2] are not credible ethical outlooks, because he thinks they are ill-suited for any of the four roles that he regards as important ones important for

[1] "Ethics Beyond Moral Theory," p. 7.
[2] Chappell actually uses the expression "moral theory," which I take to be equivalent to the expression "ethical theory" as I use it.

ethical outlooks to play.[3] I am not convinced by his arguments, but I don't have the space to address them here, so I will proceed on the assumption that mainstream ethical theories are credible ethical outlooks and leave the analysis of Chappell's arguments for another occasion.

Related to this is the important question of whether ethical outlooks that are *not* ethical theories are eligible to be included among the views that my approach takes into account. For now I will simply say that, though I cannot see *a priori* any reason why not, I think their eligibility must be determined on a case by case basis, and in order to make such determinations I would need a more fully developed account than I now possess of the criteria of viability that I will present in the next section. The development of such an account is something that I must also leave for another occasion.

3. *Criteria of Viability*

Given that there may be ethical theories that have not yet been thought of by anyone, it is not certain how many of them there are. But it is certain that not all of them are created equal. Thus, in order to carry out a project like mine, it will be necessary to develop criteria of viability that one can use to eliminate theories that are inadequate and thus narrow down the range of theories one will have to consider. In this section I will list some of these criteria.

One criterion is *cohesiveness*. Ethical theories cannot contain contradictions, but more than that, their components must be mutually supporting and fit together well. That is, an ethical theory can't just be some arbitrary set of statements about what one should do that happens to be logically consistent. Some parts of the theory must provide a *rationale* as to *why* such-and-such is right, wrong, or permissible. Furthermore, these parts and their rationales must be subsumed under some common principles, or be such that relevantly similar actions receive relevantly similar evaluations, and for relevantly similar reasons.

---

[3] Here is Chappell's characterization of these roles:

> We want our ethical outlook to be something which, in real time, can be the source of our reasons to act (*motivation*), and which can structure our thinking and deciding about how to act as it actually happens (*deliberation*). We also want our ethical outlook to be something which, offline, can articulate and deepen our understanding of what counts as good or bad and right or wrong action, and why (*explanation*); and we want it to be something which can explain what will or would be good or bad and right or wrong action, in future or hypothetical situations that we ourselves have not actually met, but which we or others might conceivably meet (*prediction*) ("Ethics Beyond Moral Theory," pp. 12-3).

Another criterion is *comprehensiveness*: An ethical theory cannot merely tell one what one should do in just a few cases. It should give one guidance that applies to a large number of cases of various kinds.

An ethical theory also has to have *verifiable* implications for one's behavior. That is, one's obligations must be such that it is in principle possible for one to discover what they are if one makes the effort to do so. If an ethical theory says, for instance, that in cases having feature *F* one should do *x* and in cases having feature *G* one should not do *x*, there must be a way for one to recognize that one is in a case which has feature *F* or a case which has feature *G*. If this were not so, one could only do *x* or fail to do *x* in the appropriate kind of case through a lucky guess. Even if one assumes that one really would have obligations in such a skeptical scenario, it would be *pointless* for one to try to *find out* what they were.

Yet another criterion derives from the old but venerable principle that "'ought' implies 'can'". Ethical theories must be *psychologically plausible*: One cannot be obligated to do something if it is psychologically impossible for one to do it. And if doing something is possible but difficult, an ethical theory which prescribes doing it is less viable the more difficult it is to do it.

As with any kind of theory, ethical theories should not contain any statements that have been shown empirically to be false. They should, in other words, be *empirically adequate*. This seems obvious enough, but if one takes one's theorizing seriously it requires that one should make the effort to see if the ethical theories that one is considering are consistent with any relevant scientific theories or bodies of knowledge.

The guidance that an ethical theory provides must also be appropriately *specific*, that is, it should not be so vague that it doesn't really recommend anything in particular.

Finally, however specific an ethical theory's guidance may be, when considering it one needs to ask oneself, "Do those who know this theory best agree about *what* it recommends and what it doesn't?" The more difficult it is for the relevant experts to agree on how to interpret the theory, the more it lacks *interpretational stability* and the less viable it is.

I will make no claim that the above list of criteria is complete, but one has to start somewhere. However, I *will* claim that these criteria are both necessary and useful for my project.

4. *A Sketch of the Solution in Action: Applying the Criteria of Viability*

I will now sketch how my project might be carried out in practice. At this point trying to use it to get *detailed* ethical guidance about some issue would be premature, because my approach is not sufficiently worked out, and trying to resolve a realistic moral dilemma satisfactorily is too large an undertaking for an incomplete sketch of an approach. What I *can* do is describe, in a somewhat general way, how one *might* apply my approach to a particular issue and to show what kinds of considerations might be relevant to it. This will help us learn how to use my approach to get detailed ethical guidance when it is finally well-developed enough to do so.

The ethical issue I will address is famine relief. Many people in developing countries have an inadequate supply of food, or are malnourished because the food they do have doesn't contain enough of the right nutrients. There is much that we in developed countries can do to help them, such as sending them food or giving them the money and resources they need to build up their economic infrastructure. I think most everyone would grant that it is *good* to contribute to famine relief, but the question that concerns us here is: Are we *obligated* to donate to this cause?

We will begin by examining some candidate ethical theories. I will consider utilitarianism, Kantianism, W. D. Ross's deontological theory, virtue ethics, John Rawls's contractarianism, and libertarianism; but because this is only an outline, they will each receive only a tiny fraction of the attention they deserve. What we have to do is to use our criteria of viability to see if we can eliminate any of these theories from consideration.

First, there is utilitarianism. Utilitarian theories hold that the ultimate goal of morality is to maximize the amount of goodness or value in the world. Some work in psychology seems to undermine the branch of utilitarianism known as "hedonic" utilitarianism, which regards the obtainment of happiness, pleasure, or other feelings of well being as the ultimate good or value that we are obliged to maximize. In his book *Ethics Done Right*, Elijah Millgram argues that psychological research shows that, "If good fortune strikes, you will be briefly elated. But as you become accustomed to your new situation, the elation will wear off, you will find new things to be dissatisfied about, and very soon you will be about as happy (or as unhappy) as before."[4]

[4] "What's the Use of Utility?", *Ethics Done Right*, p. 35

Similarly, on the same page he says that after a disaster you will eventually find pleasure in other things and come to feel about the same as you once did. "If one's utility or happiness is thought of as being a matter of how one feels, modulo short-lived fluctuations," Millgram says, "it does not look like there is much in the normal run of things that one can do to make people more or less happy."[5] Furthermore, he says,

> Evidently, utility does not covary with how well or how badly off one is. Rather, it indicates whether one's circumstances are getting, or have just gotten, better or worse. […] [Utilitarians] took it to be the goal, and the sole bearer of value. In fact, however, it plays a very different role in our mental economies. The cognitive function of utility is, I suggest, not to be, or stand in for, the absolute level of one's welfare, but to alert one to changes in it [footnotes omitted] ("What's the Use of Utility?", *Ethics Done Right*, p. 39).

If all this is true, the "pursuit of pleasure"—the constant search for a steady stream of pleasant experiences—is really a search for a will o' the wisp, and hedonic utilitarianism has presuppositions that violate the criterion of empirical adequacy, as well as the criterion of psychological plausibility. For now this leaves us with non-hedonic forms of utilitarianism, such as ideal utilitarianism and preference utilitarianism, as the only viable utilitarian theories.

According to ideal utilitarianism, the thing to be maximized is the amount of "intrinsic goodness" in the world. Pleasant experiences could be intrinsically good, but on this view other things are intrinsically good as well. The persons who have such experiences are arguably prime examples of intrinsically good things, as, arguably, are other sentient beings. Perhaps all living things, as well as some non-living ones, also count as intrinsically good. According to preference utilitarianism, on the other hand, the thing to be maximized is the satisfaction of preferences. Sentient beings, on the whole, prefer having pleasant experiences to having negative or neutral ones, but they can also prefer some things to others even when none of the alternatives would have an impact on how they feel. Furthermore, since only sentient beings can have preferences, non-sentient beings have no moral status on this view, except insofar as sentient beings have preferences concerning them. Both of these views merit further discussion, but in the interest of conserving space I will simply assume they are viable and leave their assessment for future work.

What about Kantianism, which is perhaps the most prominent example of a deontological theory? The core of Kant's theory is his Categorical Imperative. In different places he formulates it in different yet related ways, but the one that will interest us here is the Principle of

[5] "What's the Use of Utility?", *Ethics Done Right*, p. 36

Universalizability. The principle is, basically, "that one act only according to maxims of which one can at the same time will that they become universal laws."[6] As Millgram says,

> The idea is that self-frustrating plans of action are the analogs, in practical reasoning, of the kind of incoherence that contradictory beliefs amount to in theoretical reasoning. Uncontroversial models for such self-frustration can be found in means-end incoherence, as when you decide to go to New York, but tear up the ticket that would get you there… . […] You cannot coherently intend a self-frustrating plan of action, and the CI-procedure [i.e., checking one's maxims using the Principle of Universalizability] is presented as a way of checking whether you are proposing something that you cannot coherently intend. […] The question is: "*Can* you (not: *do* you) will that everybody do as you are proposing to do yourself?"("Does the Categorical Imperative Give Rise to a Contradiction in the Will?"*Ethics Done Right* p. 91).

Millgram has a very interesting objection to the CI-procedure. According to him, Kant's principle of universalizability is self-refuting. This is because "…successful agency requires exceptions from others' policies… "[7], which violates Kant's Principle of Universalizability: Universal laws just don't have exceptions. But certain practices presuppose that in some cases we will make exceptions to the rules that govern our actions, for "…one's stake in one's own agency is such that one cannot endorse having it undermined by being deprived of the exceptions that are its precondition. And if that in turn is correct, then it is forbidden to act on the CI-maxim…".[8]

One could try to handle this difficulty by building exceptions into one's maxim of action, but Millgram thinks this will not work. The following example illustrates the problem well:

> If everyone who needed to were allowed to leave his car in the lot with his blinkers on, people would soon notice this option and start planning around it. They would count on being able to dash upstairs to make a phone call, or to deliver flowers, or whatever; and so they would plan on making that phone call or delivering the flowers. The parking lot would very quickly be full of briefly parked cars. The primary users of the lot would not be able to get in and out, and the visitor who needed an exception to the rule would generally find all the free spaces already taken. That is, what I needed to make my day work was an exception, and not an adjusted rule (*Ethics Done Right* p. 104).

If Millgram is right, Kant's moral theory is undermined—at least insofar as it depends on the Principle of Universalizability—and is ruled out by the criterion of cohesiveness. Perhaps Kant's other formulations of the Categorical Imperative would fare better. If not, deontologists will have to look elsewhere for a foundation on which to build their theories.

---

[6] [footnote omitted] ("Does the Categorical Imperative Give Rise to a Contradiction in the Will?", *Ethics Done Right* p. 90).
[7] "Does the Categorical Imperative Give Rise to a Contradiction in the Will?", *Ethics Done Right* p. 97
[8] "Does the Categorical Imperative Give Rise to a Contradiction in the Will?", *Ethics Done Right* p. 97

One alternative would be the ethical theory of W. D. Ross, which is a less extreme example of a deontological theory. Although I count it as a form of deontology, it could also be classified as a form of particularism, or perhaps more accurately as occupying a no man's land between the two. In any case, Ross's view is roughly as follows. Instead of having one fundamental moral principle from which all our duties can be derived, he posits a variety of duties which he regards as mutually irreducible.[9] Ross distinguishes one's *prima facie* duties from one's duty proper, or actual duty. A *prima facie* duty is that which appears to be one's duty; it counts in favor of one's performing an action, and would be one's duty proper—what one *actually* should do—if it were the only morally significant consideration.[10] Oftentimes, however, it is *not* the only morally significant consideration. One can have multiple *prima facie* duties which conflict with each other. For example, one's *prima facie* duty not to lie can conflict with one's *prima facie* duty not to tell an inquiring murderer where one's friend is hiding. Ross also holds that some *prima facie* duties are more stringent than others, which can resolve a conflict between *prima facie* duties. In this case, one's *prima facie* duty not to lie is much less stringent than the *prima facie* duty to preserve one's friend's life, and so in this case preserving one's friend's life is one's duty proper.

One criticism of Ross's view is that it is not systematic enough. David McNaughton puts the objection like this (though he does not endorse it himself):

> Common-sense morality appeals to a large variety of moral principles, which have no discernible structure. Intuitionism [a family of views which hold that our moral beliefs are justified by our *intuitions*, i.e., what strikes us as being true, about moral matters] does not attempt to systematize ordinary morality, but simply mirrors it. An intuitionist, such as Ross, merely presents us with a more or less arbitrarily selected list of the more common (prima facie) duties, and announces them to be self-evident. Since there is no structure to this list, there seems to be no explanation of why some items are on the list and not others, and therefore no room for rational debate in the event of disagreement about what should be included. ("An Unconnected Heap of Duties?", *Ethical Theory* p. 762)

If this objection is on target, Ross's theory runs afoul of both the criterion of cohesiveness and the criterion of interpretational stability.

However, McNaughton thinks he has a reply to this objection. He says that Ross's approach *does* try to make our common-sense moral intuitions more systematic, and it does so by trying to derive all of our duties from a fairy short list of basic duties. The main difference

[9] "What Makes Right Acts Right?", *Ethical Theory* pp. 754-5
[10] "What Makes Right Acts Right?", *Ethical Theory*, p. 754

between Ross's view and utilitarian views, McNaughton thinks, is that utilitarians recognize only one basic duty—to maximize the good—while Ross recognizes several.[11] McNaughton goes on to examine what the relation of derivation might consist in, and gives a few examples of how some of our duties might be derived from our most basic ones. Thus McNaughton thinks that there is a way to tell, in principle, which duties are on the list of our most basic duties and which are not. "Critics of intuitionism," McNaughton says, "are wont to point out that different intuitionist philosophers cannot agree about which are the basic duties, as if this were itself a sufficient refutation of the theory. But this would only be an objection to intuitionism if the theory held that the contents of the list should be immediately obvious, which it does not. What is important is that there should be some rational and principled way to settle such disputes, and this is what I have tried to show."[12]

But I'm not sure that McNaughton has shown this, for two reasons. First, the examples he gives show how some of the derivations *might* go, but given the great number and variety of the duties we commonly take ourselves to have, it's not clear whether all of them could be derived from Ross's list (or something close to it) in a similar way. Second, in the more than eight decades since Ross published his book *The Right and the Good*, neither his list of basic duties nor any other has, as far as I know, gained widespread acceptance among ethicists, and if it were really possible in principle to determine which duties are on the list and which are not, it would be unlikely that the lack of consensus would have persisted for so long. Nevertheless, this is not sufficient to conclusively rule out Ross's theory, so I will not reject it as not being viable. Instead, I will regard it as being *doubtful*, though still worthy of consideration.

Let us move on to virtue ethics. In virtue ethics the main focus is, unsurprisingly, on living a virtuous life. To do the right thing is, roughly, to act as an ideally virtuous person, or at least a sufficiently virtuous person, would act themselves. In order to become a virtuous person, it is necessary to possess character traits, or to develop habits, that lead one to act in a virtuous manner. But "Virtue ethicists also claim that having a virtue, which is something that comes by degrees, contributes to making one's life a good one—to what Aristotle called *eudaimonia*, or flourishing. A life that exhibits the virtues is for that very reason a better life… ."[13]

[11] "An Unconnected Heap of Duties?", *Ethical Theory* p. 762
[12] "An Unconnected Heap of Duties?", *Ethical Theory* p. 762
[13] *Experiments in Ethics*, p. 35

Virtue ethics also has problems regarding psychology. Kwame Anthony Appiah, in *Experiments in Ethics*, presents research in social psychology that seems to show that people do not have stable character traits of the sort virtue ethicists suppose them to have. For example, Appiah says,

> In the past thirty years or so, broader psychological evidence against globalism [the belief that people have "…consistent dispositions to respond across contexts under the guidance of a certain value…".[14]] has been accumulating. Back in 1972, Alice M. Isen and Paul Levin found that when you dropped your papers outside a phone booth in a shopping mall, you were far more likely to be helped by people if they had just had the good fortune of finding a dime in the phone's coin-return slot. A year later, John Darley and Daniel Batson discovered that Princeton seminary students, even those who had just been reflecting on the Gospel account of the Good Samaritan, were much less likely to help someone "slumped in a doorway, apparently in some sort of distress," if they'd been told that they were late for an appointment. In a 1975 study, people were much less likely to help someone who "accidentally" dropped a pile of papers when the ambient noise level was 85 decibels than when it was 65 decibels. More recently, Robert Baron and Jill Thomley showed that you were more likely to get change for a dollar outside a fragrant bakery shop than standing near a "neutral-smelling dry-goods store" [footnote omitted] (*Experiments in Ethics*, pp. 40-1).

The problem that psychology poses for virtue ethics is more acute than the problem it poses for utilitarianism. If hedonic utilitarianism fails, there are still ideal utilitarianism, preference utilitarianism, and other, related forms of consequentialism to fall back on. But if the research Appiah cites is correct *there are no virtuous habits or character traits for people to cultivate*, and hence all forms of virtue ethics are threatened. Of course, there is more to be said by those on both sides of this debate, but for the purposes of this essay I will set virtue ethics aside.

The next theory up for consideration is John Rawls's version of contractarianism. In general, contractarianism is the view that the demands of morality are ultimately grounded in an agreement or contract, whether hypothetical or real, between persons who, though perhaps not selfish, are at least self-interested, and who sacrifice some of the rights or freedoms they would enjoy as solitary individuals in the "state of nature" for the sake of the greater benefits they would obtain from living with others as part of a society. Rawls offers his theory as one of justice and related notions, as opposed to a theory of morality generally,[15] but the transition from one to the other is natural and for our purposes we can ignore the distinction between them. For Rawls, the principles of justice are the terms of an agreement—his version of the social contract—reached by persons in "the original position"—his version of the state of nature. According to him,

---

[14] John M. Doris, *Lack of Character*, pp. 61-2; quoted in *Experiments in Ethics*, p. 38
[15] "A Theory of Justice," *Ethical Theory*, p. 631-4

> [The original position] is understood as a purely hypothetical situation characterized so as to lead to a certain conception of justice. Among the essential features of this situation is that no one knows his place in society, his class position or social status, nor does any one [sic.] know his fortune in the distribution of natural assets and abilities, his intelligence, strength, and the like. I shall even assume that the parties do not know their conceptions of the good or their special psychological propensities. The principles of justice are chosen behind a veil of ignorance. This ensures that no one is advantaged or disadvantaged in the choice of principles by the outcome of natural chance or the contingency of social circumstances. Since all are similarly situated, and no one is able to design principles to favor his particular condition, the principles of justice are the result of a fair agreement or bargain. [footnote omitted] ("A Theory of Justice," *Ethical Theory*, pp. 631-2)

I find less to criticize in Rawls's theory than in the previous ones because as it stands it seems to be consistent with almost all of the criteria of viability. For example, it seems to be cohesive: There appear to be no obvious inconsistencies, and the benefits that the persons who participate in the original position would obtain by adhering to the terms of their agreement provide a good rationale for acting as the agreement prescribes. Furthermore, since the original position treats all of the participants symmetrically, its prescriptions and prohibitions are interpretationally stable: Different persons who imagine themselves to be in the original position and who deliberate in accordance with its strictures should reach the same conclusions about what they should or shouldn't do. Rawls's theory also seems to be empirically adequate, because the original position is itself merely hypothetical, and it can't conflict with experience if it has no empirical presuppositions. The only empirically significant assumption it makes is that it is *possible* for people to reason as the theory says they should. Regarding the other criteria, the only potential problem that I can see is that Rawls's theory might be inconsistent with the criterion of comprehensiveness, because its purpose is to be a theory of justice rather than a theory of morality in general. I honestly don't know whether it could be successfully developed into a comprehensive ethical theory or not. But in this case it might not matter. As I will argue in the next section, one can give a Rawlsean argument in favor of donating to famine relief, so even if Rawls's theory has limitations as an ethical theory, it can still be used to get ethical guidance in this instance.

Finally, we come to libertarianism. Libertarianism can be taken as a political view, an ethical view, or perhaps both. Since the purpose of my project is to get ethical guidance, I will only consider libertarianism as an ethical view. According to Engel, "Strict libertarians insist that although we have negative duties to do no harm, we have no positive duties to assist others. […] Because they deny the existence of positive duties, libertarians also contend that it would *not* be

wrong of you *not* to save the lives of numerous starving children by sending a modest portion of your income to famine-relief organizations."[16] I'm not sure how many libertarians of either the political or moral variety would self-identify as being "strict libertarians" in this sense, so to avoid attributing views to them that they don't really hold I will stipulate that in the ensuing discussion the term 'libertarian' will be *definitionally equivalent* to the term 'strict libertarian' as Engel uses it. I will also assume that infringing on others' liberties counts as "harm" on this view.

Is libertarianism, so construed, consistent with our criteria of viability? I think it is, if only because of its stark simplicity. It is *cohesive*, because similar actions receive similar evaluations in similar situations: Cases where an action would harm others receive similar evaluations, because one is forbidden from doing it. Cases where an action would not harm others receive similar evaluations, because one is permitted to do it or not to do it as one sees fit. It is (vacuously) *comprehensive*, because it tells one what one should do a large number of diverse cases: In every case where one's action would harm another, one cannot do it; in all other cases, of whatever sort, one can do whatever one likes. It has *verifiable* implications for one's behavior, because it is usually fairly easy to tell when a proposed course of action would be likely to harm others and when it would not. It is also *psychologically plausible*: It places very few constraints on one's behavior, so it is extremely easy to comply with. Granted, most people feel an impulse to help others in need, but helping them is not *incompatible* with libertarianism, it's just not *required* by it. It makes no significant empirical assumptions, so it is *empirically adequate*. Its guidance is *specific*: The requirement that one is forbidden to harm others is not significantly vague, nor is the corresponding permission to do what one feels like in other cases. Finally, its guidance is *interpretationally stable*: Since it is usually fairly easy to tell when a proposed course of action would be likely to harm others and when it would not, there is not likely to be much disagreement about when one is forbidden from doing something and when one is not.

[16] "Hunger, Duty and Ecology," *Environmental Ethics*, p.472

5. *A Sketch of the Solution in Action: Comparing Ethical Theories*

Having (slightly) whittled down the number of ethical theories we will have to consider, we are finally in a position to see if we can reach a verdict on the issue of donating to famine relief, and if we can, to determine what that verdict is. The first argument for the obligatoriness of donating to fame relief is Peter Singer's utilitarian "Basic Argument," taken from his book *The Life You Can Save*:

> First Premise: Suffering and death from lack of food, shelter, and medical care are bad.
> Second Premise: If it is in your power to prevent something bad from happening, without sacrificing anything nearly as important, it is wrong not to do so.
> Third Premise: By donating to aid agencies, you can prevent suffering from lack of food, shelter, and medical care, without sacrificing anything nearly as important.
> Conclusion: Therefore, if you do not donate to aid agencies, you are doing something wrong. (*The Life You Can Save*, pp. 15-6)

This argument is clearly valid, but are its premises true? If utilitarianism is true, I think they are. Now, as stated, this argument may seem to require the truth of *hedonic* utilitarianism, but in reality it doesn't. Though ideal utilitarians would not regard suffering as the *only* evil, they would, I think, still regard it as *an* evil that we should try very hard to prevent. Similarly, while preference utilitarians think that we have preferences that concern things other than pleasure and pain, that does nothing to detract from the fact that we *do* have such preferences and that people strongly prefer not to suffer. Thus both of the remaining forms of utilitarianism would require us to donate to famine relief.

I think there is also a Kantian argument in favor of contributing to famine relief, one which is not ruled out by what was said in section 4. As I mentioned there, other formulations of the Categorical Imperative might fare better than the Principle of Universalizability. Onora O' Neill uses one such formulation of Kant's Categorical Imperative, the Formula of the End in Itself, to argue for contributing to famine relief. This says that one should "Act in such a way that you always treat humanity, whether in your own person or in the person of any other, never simply as a means but always at the same time as an end."[17] Given this imperative, one could make the following argument for contributing to famine relief:

> Since finite rational beings cannot generally achieve their aims without some help and support from others, a general refusal to help and support amounts to failure to treat others as rational and autonomous beings, that is, as ends in themselves. […] Since hunger, great poverty, and powerlessness all undercut the possibility of autonomous action, and the requirement of treating others as ends in themselves demands that Kantians standardly act to support the possibility of

[17] "Kantian Approaches to Some Famine Problems," *Ethical Theory*, p. 554

> autonomous action where it is most vulnerable, Kantians are required to do what they can to avert, reduce and remedy hunger. They cannot of course do everything to avert hunger: but they may not do nothing. ("Kantian Approaches to Some Famine Problems," *Ethical Theory*, p. 556)

Thus, if the Formula of the End in Itself can withstand scrutiny, we will have at least one Kantian argument for contributing to famine relief.

In section 4 I regarded Ross's account as doubtful, though I felt that the objection was not strong enough to rule it out completely. How, then, should we proceed? In such a case, I think one should assume that the doubtful theory is viable and see if one can use it to assess the moral status of the action one is considering. If one can, and it turns out that the theory is viable after all, well and good. But because it might also turn out not to be viable, we should check to see if there is sufficient agreement among the reaming theories to reach a verdict about the moral status of the action under consideration without appealing to the doubtful theory. If the verdict doesn't change when we drop the assumption that that theory is viable, it really doesn't matter whether the theory is viable or not, for the verdict is the same in either case. So if both the assumption that a doubtful theory is viable and the alternative assumption that it is not viable lead to the same verdict when the set of other theories that are being compared is held constant, we really don't need to know whether the doubtful theory is viable.

That being so, I propose the following Rossian argument for contributing to famine relief: Ross acknowledges that persons have a duty of beneficence, which is, barring possible overriding factors, to bring about as much good as we can.[18] In poor countries, lack of food and malnourishment lead to extreme suffering. So given that—what I take to be a very safe assumption—relieving suffering is good, it follows from Ross's view that we should donate what we can to famine relief as long as our doing so doesn't conflict with any stronger duties that we have. I think that there are very few such duties, and if I'm right about that that means that Ross's theory entails that we should do a lot to relieve or prevent famine, in all probability a lot more than most of us actually do.

The fourth argument in favor of donating to famine relieve is based on Rawls's contractarian theory. Someone who accepts Rawls's theory could argue as follows. If the persons in the original position are under a "veil of ignorance," they should not be taken to know either

[18] "…if we are ever under no special obligation such as that of fidelity to a promisee or of gratitude to a benefactor, we ought to do what will produce most good; and that even when we are under a special obligation the tendency of acts to promote general good is one of the main factors in determining whether they are right" ("What Makes Right Acts Right?", *Ethical Theory*, pp. 758-9).

which society they will be a part of, or which of their fellows will belong to the same society as them. This is because these restrictions will ensure, just like the restrictions that Rawls himself imposes, "…that no one is advantaged or disadvantaged in the choice of principles by the outcome of natural chance or the contingency of social circumstances."[19] Thus these persons should—if only out of prudence—construct a system of government such that persons in different societies would offer each other aid if there is a genuine need for it, provided that it wouldn't have a significant negative impact on the donor society. For if one is in the original position, and hence doesn't know which society one will belong to, one doesn't know whether one will belong to a recipient society or a donor society (or perhaps neither). If one ends up in a recipient society the need for assistance could be very great, and if one ends up in a corresponding donor society the cost will probably be minor in comparison.

The fifth argument for donating to famine relief comes from an unlikely source: libertarianism. The argument can be derived from Singer's evaluation of this theory. Singer, like Engel, is no friend of libertarianism. He notes that, if the libertarian view were correct, we would be required to get rid of state supported welfare schemes and state funded healthcare.[20] But leaving that aside, libertarianism may not be as irrelevant to famine relief as it seems. Ironically, it turns out that the libertarian argument for donating to famine relief can be derived from a libertarian objection to the obligatoriness of such donations. Jan Narveson gives voice to such an objection when he says,

> I will take it as given that we are certainly responsible for evils we inflict on others, no matter where, and that we owe those people compensation. Not all similarly agree that it is not in general our duty to make other people better off, and therefore not in general our fault when people are not better off than they happen to be, even if perhaps we *could have* made them so by efforts of our own. Nevertheless, I have seen no plausible argument that we *owe* something, as a matter of general duty, to those to whom we have done nothing wrong. ("We Don't Owe Them a Thing!", p. 1)

In examining this passage, Singer argues that if we have harmed the world's poorest by being a partial cause of their poverty, then by libertarian standards we should compensate them. He has little trouble showing that we have harmed the world's poorest in that way. Such harms include depriving them of jobs and a main source of protein through overfishing, buying products from corporations who deal with corrupt dictators, and causing global warming.[21] In the last instance, global warming will cause precipitation to "…decrease nearer the equator and increase nearer the

[19] "A Theory of Justice," *Ethical Theory*, pp. 631-2.
[20] *The Life You Can Save*, pp. 28-9
[21] *The Life You Can Save*, pp. 29-33

poles. In any case, the rainfall upon which hundreds of millions rely to grow their food will become less reliable."[22] Also, as sea levels rise, densely settled delta regions and some small Pacific Island nations will be inundated or even submerged. Singer concludes—quite rightly, in my opinion—that "If we accept that those who harm others must compensate them, we cannot deny that the industrialized nations owe compensation to many of the world's poorest people."[23]

Now, one might object that most of the people in developed countries have not personally done any of these things; it is rather their governments or their corporations who have done them. Not everyone in these countries has voted for or participated in the responsible governments or has dealt with the responsible corporations, and even those who have may not have known about these activities, or had little choice about their involvement. So most of us are not personally responsible for the harms done to the people of poor countries, and hence most of us don't personally owe them any compensation.

My reply is that one might owe someone compensation for doing (or failing to do) something even if one hasn't wronged them in doing it. For example, one might justifiably steal food to feed one's family if one has no alternatives. In this case one has done nothing wrong, but in spite of that one's act of theft was non-ideal; it is *unfortunate* that one had to steal the food. I think that if one later becomes better off and is able to compensate those from whom one stole, one is obligated to do so, even by the standards of strict libertarianism. Even if those of us who live in developed countries have, personally, not *wronged* those who live in poor countries, most of us have still personally *benefited* from the harm that the responsible governments or corporations have done to them. Maybe most of us can't avoid doing so, but as long as we are in a position to compensate them for these harms—and most of us surely are—I think we are personally obligated to compensate them.

This case illustrates the distinction between what I will, drawing on Kant, call the *categorical content* and the *hypothetical content* of ethical theories. The categorical content of an ethical theory concerns what you should or shouldn't do no matter what, while its hypothetical content concerns what you should or shouldn't do *if* certain conditions are met. Regarding the above case, sending money to the world's poorest people *if we have wronged them* is part of the hypothetical content of libertarianism, but sending them money *simpliciter* is not a part of

[22] *The Life You Can Save*, p. 32
[23] *The Life You Can Save*, p. 33

libertarianism's categorical content, for according to libertarianism we have no positive duties. By contrast, sending money to the world's poorest people *if we have wronged them* is part of the hypothetical content of Kantianism, and sending them money *simpliciter* is a part of its categorical content. Libertarianism, then, unlike Kantianism, does not directly imply that one should donate to famine relief. But we are not therefore free to set libertarianism aside. Libertarianism is still relevant to our deliberations concerning fame relief because we *have in fact* harmed the world's poorest people. This example teaches us that, when comparing ethical theories, we should consider both their categorical and their hypothetical content.

At long last, we have our verdict: All of the viable ethical theories that we have examined prescribe that we are obligated to contribute to relieving famine. Even if we exclude Ross's theory, all of the remaining theories still agree that we must do so. As long as one of these theories is true, it follows that we should do what they all prescribe. So there is at least one issue concerning which we have found out how to know what should be so.

6. *Objections and Replies*

Having sufficiently explained my approach for getting ethical guidance, my criteria of viability, and how my project might be carried out, I would now like to address a few objections that could be raised against them. First, one could object that one must often make a snap decision about what one should do and has precious little time to compare various ethical theories. My approach could not provide guidance in such circumstances. So in what circumstances *would* it be useful?

In response, I concede that my approach will not provide someone with guidance if they encounter a novel moral dilemma and have a short time to decide what to do. However, I certainly think my approach *will* be useful in situations where one faces a moral dilemma and has a reasonable amount of time to decide what to do. For example, if you are considering whether or how much you should donate to research investigating possible treatments for cancer, it probably won't make too much of a difference whether you reach a decision now or in two months. In this case one has the time to compare different ethical theories and see if one can find enough common ground to reach a decision. But I think my approach could also be useful even when you must make a snap decision, provided that you develop the habit of pausing from time to time to reflect on what you ought to do *if* you were to face various moral dilemmas. For those

who recognize that they will often face snap decisions regarding what they ought to do, the wisdom of trying to think through such decisions beforehand will be apparent.

The next objection is similar to the first. My approach would require a fairly detailed knowledge of various ethical theories in order to get useful ethical advice, for without it the ethical theories that one is aware of might not constitute a representative sample of all the viable ethical theories there are. And even if one is aware of a representative sample, one's knowledge of some of the theories might not be detailed enough for one to be able to determine what they would tell one to do in a particular situation.

In reply, I concede that it would be difficult for a single, isolated individual to use my approach to get useful ethical guidance. But if someone is either a part of, or has access to, a group of individuals which *does* have detailed knowledge of various ethical theories—or better still, a group of individuals which is dedicated to carrying out my project—one could acquire some of their knowledge for oneself or ask them for their advice. If someone who is in a moral quandary is moral enough to be motivated to earnestly seek out ethical guidance and has the time to do so, they will probably be able find some such group. Thus I do not think that an initial lack of a detailed knowledge of ethical theories is a serious problem for such a person. It is no more of an objection against my project that it could not be practiced by an individual working in isolation than it is an objection against science that it could not be practiced by an individual working in isolation.

Third, there is *moral particularism*, which holds that the moral status of an action depends primarily on the details of the particular situation in which it takes place. This might seem to call into question the very need for ethical theories, and without them my project couldn't even get off the ground. One question is thus whether my criterion of cohesiveness rules out moral particularism. If it does, wouldn't it be a mark against my approach that it disqualifies a major ethical view? And in any case, since my approach will admittedly not yield an ethical theory, wouldn't the guidance one would get from my approach be very similar to the guidance one would get from particularism?

In response to the first question, I will give an appropriately particularist answer: It depends. There are a few different possible moral particularisms, and the criterion of cohesiveness may rule out some of them but not others. For instance, someone might embrace what I'll call an *antigeneralistic* particularism, according to which there aren't any (or are very

few) *general* moral principles which hold across different circumstances and which determine that the same types of action have the same moral status in situations that are sufficiently similar. Jonathan Dancy espouses a view of this sort:

> One of the first things that one learns when studying how reasons behave is that a certain theory, atomism, is false. Atomism holds that any feature that is a reason in favor of an action in one case will always be a reason in favor of that action wherever it occurs. The same feature always makes the same reason; or, a reason is a general reason. This theory is false; something that is a reason in favor of an action in one case may in another case be no reason at all, or even a reason against action. It all depends on the circumstances; reasons are sensitive to context. ("An Unprincipled Morality," *Ethical Theory*, p. 771)

And shortly afterwards, Dancy says:

> …most people think that moral reasons are based on principles. But, as I will argue, if atomism is false, there can be no moral principles. *Moral particularism* holds that, because of the falsehood of atomism, there are plenty of moral reasons but no moral principles. ("An Unprincipled Morality," *Ethical Theory*, p. 772)

If this is correct, one could not come to know what the moral status of an action is by inferring it from one's knowledge that it is an instance of a type of action that is inherently associated with a certain moral status. One would instead have to come to know the moral status of an action by means of one's perception of the morally relevant features of a specific situation. If we understand an ethical theory as something which involves subsuming our evaluations of actions "under some common principles," or which is "such that relevantly similar actions receive relevantly similar evaluations, and for relevantly similar reasons," then according to antigeneralistic particularism there could be no *generalistic* ethical theory[24] as to why particular actions have the moral status they do. For the same morally relevant features that make an action right in one situation could in other situations make that action wrong, and in still other situations have no relevance to its moral status. So different antigeneralistic particularists would seem to have no theoretical means to resolve possible disputes about the moral status of particular actions, and would thus seem to be unable to give others useful ethical advice. If this is right then the criterion of cohesiveness, as well as the criterion of interpretational stability, would require us to reject antigeneralistic particularism because it could not be used to give us guidance even if it were true.

However, there are other particularist views, such as those of Margret Little and Gerald Dworkin, which may be more congenial to generalistic ethical theories. On Little's view one *can* make ethical generalizations, but they are not statistical generalizations. Rather, they mark

[24] I.e., an ethical theory which holds that there are general moral principles.

certain conditions as being normal or privileged, in the sense that they distinctively reveal something's nature. "On this theory," she says,

> [W]hen we say "ceteris paribus, lying has a negative valence," we are not saying that it always, must, or even usually has that status; we are saying that this is the valence it has in conditions that are privileged in various ways. When we say "all things equal, pain is bad-making," we are not saying that pain always carries this valence, nor merely asserting that it usually does in our neck of the woods. We are saying, instead, that pain is defeasibly bad-making; it has a default negative valence. Where lying and pain lack this valence, as they sometimes or even often do, it is because they occupy a context defective by morality's own lights, or again because the cases in question are operating as variations that cannot be understood except by reference to a paradigm that carries the privileged valence. ("On Knowing the "Why"," *Ethical Theory*, p. 782)

As long as one has some way of knowing what the privileged valences are, one can give similar actions similar evaluations in circumstances where the appropriate valences are privileged. This view, then, appears to be compatible with my approach.

In his paper Dworkin presents psychological research which he takes to show that people do not arrive at moral judgments about particular cases by subsuming them under general principles.[25] He suggests a model according to which people compare novel cases to paradigm cases. Paradigm cases are particular situations in which a moral decision was made, and for which the decision is regarded as having been a good one. If a novel case is sufficiently similar to the paradigm, a person makes the same moral judgment about it as the one they made about the paradigm case.[26] Provided that different moral agents have access the same paradigms and largely agree on how similar different cases are to each other, this view would also appear to be compatible with my approach.

In response to the second question, I agree that ruling out particularism would be a mark against my approach if there were only one version of particularism and it was well supported. But if I'm right that there are different versions, like those of Little and Dworkin, which are compatible with my approach, there is no problem. And if I'm wrong, and an approach like Dancy's is the only coherent version of particularism, that would still only tell against my view if it was well supported, and I think there's room for debate about that.

Finally, as for the issue of whether my approach would give guidance which is similar to the guidance one would get from particularism, I must concede that the guidance itself might be similar in the sense that both on my approach and on particularism there is no single overarching ethical theory from which one could derive verdicts concerning particular cases. Yet there are

[25] "Unprincipled Ethics," *Ethical Theory*, pp 790-5.
[26] "Unprincipled Ethics," *Ethical Theory*, p 793.

two important differences. First, on my approach there would be a straightforward justification for the verdict one would get in every particular case where one would get a verdict at all, because all or most ethical theories would agree about the moral status of the action one is considering. On an antigeneralistic particularist approach, one would have to rely on an alleged ability to *just perceive* that some features of that circumstance are morally relevant and have the moral valence they do even though they might be morally irrelevant or have a different moral valence in other circumstances. On particularist views like Little's or Dworkin's one wouldn't have to rely on this, but these views are among those that my approach could take into account, and so they wouldn't be in competition with it. Second, due to the fact that on my approach all or most ethical theories would agree about the moral status of the action one is considering, different moral agents who use my approach will get the same verdicts about the same actions in the same circumstances, as long as they take into account the same ethical theories and use the same criteria of viability. So while one cannot use my approach to develop an ethical theory, it will nevertheless give one guidance that is interpretationally stable.

7. *Conclusion*

In this essay I have only been able to give a very rough sketch of my approach. In the future I hope to work it out in detail and apply it to actual cases of moral decision making. If it works then there is a way to get ethical guidance which bypasses the debate over which ethical theory is true. In spite of that, my approach is one in which ethical theories play an essential part. Thus ethical theorizing can be both helpful and relevant to one's moral decision making, though perhaps not in the way one might initially think.

References

Appiah, Kwame Anthony. *Experiments in Ethics*. Harvard University Press, 2008

Chappell, Timothy. "Ethics Beyond Moral Theory". *Philosophical Investigations* 32 (3):206-243. 2009

Dancy, Jonathan. "An Unprincipled Morality". *Ethical Theory: An Anthology*, edited by Russ Shafer-Landau. Blackwell Publishing Ltd, 2007

Doris, John M. *Lack of Character: Personality and Moral Behavior*. Cambridge University Press, 2002

Dworkin, Gerald. "Unprincipled Ethics". *Ethical Theory: An Anthology*, edited by Russ Shafer-Landau. Blackwell Publishing Ltd, 2007

Engel, Mylan (Jr). "Hunger, Duty, and Ecology: On What We Owe Starving Humans". *Environmental Ethics*, Edited by Louis P. Pojman and Paul Pojman. Thomson, Wadsworth, 2008

Little, Margret Olivia. "On Knowing the "Why": Particularism and Moral Theory". *Ethical Theory: An Anthology*, edited by Russ Shafer-Landau. Blackwell Publishing Ltd, 2007

McNaughton, David. "An Unconnected Heap of Duties?" *Ethical Theory: An Anthology*, edited by Russ Shafer-Landau. Blackwell Publishing Ltd, 2007

Millgram, Elijah. *Ethics Done Right*: *Practical Reasoning as a Foundation for Moral Theory*. Cambridge University Press, 2005

Narveson, Jan. "We Don't Owe Them a Thing! A Tough-minded but Soft-hearted View of Aid to the Faraway Needy". *The Monist*, vol. 86, no. 3, pp. 419-433. 2003

O' Neill, Onora. "Kantian Approaches to Some Famine Problems". *Ethical Theory: An Anthology*, edited by Russ Shafer-Landau. Blackwell Publishing Ltd, 2007

Pojman, Louis P. and Paul Pojman (eds.), *Environmental Ethics: Readings in Theory and Application, Fifth Edition*. Thomson Wadsworth, 2008

Ross, W. D. "What Makes Right Acts Right?" *Ethical Theory: An Anthology*, edited by Russ Shafer-Landau. Blackwell Publishing Ltd, 2007

Shafer-Landau, Russ (Editor), *Ethical Theory: An Anthology*. Blackwell Publishing Ltd, 2007

Singer, Peter. *The Life You Can Save: Acting Now to End World Poverty.* Random House, New York, 2009