LES ATELIERS DE L'ÉTHIQUE

LA REVUE DU CREUM

UNE REVUE MULTI—
DISCIPLINAIRE SUR LES
ENJEUX NORMATIFS DES
POLITIQUES PUBLIQUES ET
DES PRATIQUES SOCIALES.

2

VOLUME 3 NUMÉRO 1
PRINTEMPS/SPRING 2008

A MULTIDISCIPLINARY
JOURNAL ON THE
NORMATIVE CHALLENGES
OF PUBLIC POLICIES
AND SOCIAL PRACTICES.

NOTE AUX AUTEURS

Un article doit compter de 10 à 20 pages environ, simple interligne (Times New Roman 12). Les notes doivent être placées en fin de texte. L'article doit inclure un résumé d'au plus 200 mots en français et en anglais. Les articles seront évalués de manière anonyme par deux pairs du comité éditorial.

Les consignes aux auteurs se retrouvent sur le site de la revue (www.creum.umontreal.ca). Tout article ne s'y conformant pas sera automatiquement refusé.

GUIDELINES FOR AUTHORS

Papers should be between 10 and 20 pages, single spaced (Times New Roman 12). Notes should be placed at the end of the text. An abstract in English and French of no more than 200 words must be inserted at the beginning of the text. Articles are anonymously peer-reviewed by members of the editorial committee.

Instructions to authors are available on the journal website (www.creum.umontreal.ca). Papers not following these will be automatically rejected.

61

VOLUME 3 NUMÉRO 1
PRINTEMPS/SPRING 2008
ARTICLES:

FRIENDSHIP AND THE GROUNDS OF REASONS

DIANE JESKE
UNIVERSITY OF IOWA

RÉSUMÉ

L'amitié de même que d'autres relations intimes ont créé des difficultés pour les philosophes moraux. Bien que la morale semble exiger que nous demeurions impartiaux, l'amitié semble donner naissance à des obligations de partialité envers nos proches. Mais cette difficulté peut disparaître une fois que l'on cesse de se concentrer sur la catégorisation des raisons en tant que morales ou non morales. Cette tendance à classer les raisons comme morales vs. non morales nous mène à accorder la place d'honneur à l'étiquette « moral » et à supposer que cette catégorie est uniforme. Si nous abandonnons ces suppositions, alors les raisons données par nos proches ou nos amis ne nous sembleront plus problématiques. Il sera alors possible de voir que toutes les raisons sont, en bout de ligne, egocentriques d'une importante façon, et que la délibération doit toujours provenir d'une perspective egocentrique.

ABSTRACT

Friendship and other intimate relationships have created difficulties for moral philosophers. While morality seems to require us to remain impartial between persons, friendship seems to generate demands or obligations of partiality toward our intimates. But the difficulty can be removed once we cease to focus on categorizing reasons as moral or non-moral. This tendency to divide reasons into categories of moral vs. non-moral leads us to give those that we label 'moral' pride of place and to assume that the category must be uniform. If we abandon these assumptions, then reasons of intimacy or friendship will no longer be so puzzling. We will then be able to see that all reasons, in the end, are importantly egocentric, and that deliberation must always proceed from an egocentric perspective.

Friendships, and other intimate relationships, have often created difficulties for moral philosophers. From the perspective of someone party to such a relationship, it seems as though one is obligated to care for one's friend (or other intimate[1]) before and to a greater extent than one is obligated to care for a stranger or mere acquaintance. If one is standing on a sinking ship and can only save a friend or a stranger, flipping a coin as a procedure for deciding whom to save seems not only absurd but also repugnant. It seems that if one engages in such a procedure or ends up saving the stranger, the friend is entitled to feel outrage, hurt, anger, and betrayal.

But moral philosophers have been inclined to run into trouble when they come to justifying the attitude and resultant actions of real-world friends. In so far as we understand ourselves as being *obligated* or *required* to care for our friends out of proportion to their need, it seems that we take friendship as generating *moral demands* on us. Does morality, though, really allow us to differentiate between persons merely on the basis of their contingent relationships to us? Further, how could morality actually *demand* that we give our friends and intimates such special consideration?

It is evident then that we need to understand not only friendship but also the nature of morality if we are to resolve these tensions. In what follows, I will argue that placing too much significance on a reason's being a *moral* reason can make matters seem much more complicated and problematic than they really are. Once we start dividing reasons into categories, we tend to give those we label 'moral' pride of place and also assume that the category must be uniform. If we abandon these assumptions, reasons of friendship (or, intimacy) will no longer be so puzzling.

I will take as my focus consequentialism, and I will do so for several reasons: it is paradigmatically a moral theory, it displays some features that it has in common with other paradigmatic moral theories such as Kantianism, it has achieved (in combination with Kantianism) a kind of dominance in contemporary ethical theorizing, it is notorious for the way that it deals with intimate relationships such as friendship, and, perhaps most importantly, it is, I think, highly plausible and attractive. At the end of the paper, I will make some brief remarks about virtue theory and friendship.

UTILITARIANISM[2] AS PARADIGM

Utilitarians understand morality as involving one, and only one, fundamental principle, the principle of utility: right action is that action which, out of all alternatives available to the agent, produces the greatest net sum of intrinsic value for all sentient beings affected in the long run.[3] There are three features of this principle, as I am understanding it, that are particularly relevant for our purposes: it is universalistic, it is impartial, and it involves a non-relativistic conception of intrinsic value.

First, the principle is universalistic: it requires the agent to consider the consequences of her actions for all persons affected. Thus, in deciding how to distribute my disposable income, I am not to think only about my friends and other loved ones, but about all persons whose welfare might be promoted through use of that income.[4] Second, in considering all such persons, I am not allowed to give any greater weight to the interests of my friends; in other words, I am required to be impartial between potential beneficiaries of my actions.[5]

The universalistic impartiality of utilitarianism is in stark contrast to the commonsense thinking that I outlined in the opening to this paper. Commonsense seems to require me *not* to give equal weight to the interests of strangers and to the interests of friends. As I said, anyone who does so seems to deserve *moral* condemnation. If I am in a lifeboat with room for only one more and I pull a stranger out of the water rather than my best friend, it seems that I have not done what I ought to have done.[6]

Utilitarians, of course, have spent a lot of time and effort trying to accommodate commonsense thinking on points such as these. As Sidgwick famously argued, we are in a better epistemic and causal position for benefiting intimates such as friends than we are for benefiting strangers: we know the needs and desires of our friends and are well placed to satisfy those desires and to meet those needs.[7] But it is important to see that whatever partiality to friends utilita-

rianism can get us, it is necessarily derivative. The fundamental utilitarian principle is impartial between persons, and any partiality must be justified as being ultimately productive of maximal value for all persons considered equally. And there is certainly no guarantee that such derivative partiality will extend to pulling my best friend rather than the stranger into the lifeboat.

Now, of course, counterexamples have been the staple recourse of all those who have wanted to reject utilitarianism. Surely we ought not hang the innocent man! Surely we ought not break the death-bed promise! Surely we ought not cut people up and distribute their organs! Surely we ought to pull our best friends into lifeboats! But we have to give utilitarianism its due, and it is precisely those features of the theory that lead to counterintuitive results that make it a paradigm of a moral theory. Morality, whatever it is, has to be contrasted with self-interest, it seems, and has to involve some sort of recognition of the equal worth of all persons. How can contingencies such as friendships affect life and death moral decisions? After all, we become friends with others for any number of arbitrary reasons: she lived next door to me, he entered graduate school in the same year as I, we shared a cell in the county lock-up, etc. If all persons are of equal moral worth, how could I be justified in showing differential concern for someone because she is my friend, given that her becoming my friend might be the result of some quirk of fate?

For this same kind of reason, classical utilitarians eschewed any appeal to a relativized theory of intrinsic value. If the value of something for an agent was a matter of the agent's subjective attitude toward that thing, then the utilitarian could easily accommodate the intuitions about friendship: my caring about my friend makes my friend's life much more subjectively valuable to me than the life of the stranger.[8] But morality, it might be said, must contrast with self-interest, and an appeal to a relativized conception of intrinsic value, according to some, would turn utilitarianism into a theory about the latter rather than about the former.[9]

There is, however, another route to making utilitarianism compatible with our commonsense understanding of the demands of friendship. One could adopt utilitarianism as a moral theory, leaving its universalistic impartiality and non-relativistic theory of intrinsic value intact, but claim that moral reasons must compete against other reasons, such as reasons of self-interest and reasons of friendship. The obstacle to taking this course is another feature that many attribute to those reasons characterized as moral: they are overriding. So, it is said, even if we have reasons that are non-moral, it is rational or permissible to act on them only when there are no moral reasons supporting an opposing course of conduct. Whatever moral reasons are, some philosophers argue, they take precedence over mere self-interest, etc. For both the classical utilitarians and for Kant, rationality and morality do not come apart. (Below, I will explain why I do not accept these claims.)

THE COMMONSENSE CONTRAST

Our ordinary practices of justification, on the other hand, allow diverse kinds of facts to provide grounds for reasons. Most importantly for our purposes, the fact that Tracy is my friend seems to ground a reason for me to care for her before and to a greater extent than I have reason to care for strangers or mere acquaintances. These non-derivative reasons to be partial are, as we have seen, quite ruled out by our paradigm of a moral theory, utilitarianism. So how are we to accommodate such reasons?

The peculiar difficulty posed by reasons of friendship (or, more generally, intimacy) is that they seem to occupy a ground somewhere between the paradigmatic moral reasons and reasons of self-interest. Unlike the paradigmatic moral reasons defined by the principle of utility, reasons of friendship are partial and non-universalistic: they require the agent to take into account the interests only of those persons who happen to stand in the friendship relation to her, and they require her to give the interests of such persons greater weight in her deliberations about how to use her resources. Unlike reasons of self-interest, reasons of friendship often require sacrifices on the part of the agent: the agent is sometimes required to promote her friend's good at the cost of her own good.[10]

If we continue to think in terms of the paradigm of moral reasons as defined by utilitarianism, then, in so far as they are not impartial or universalistic, reasons of friendship seem not to qualify as moral reasons. And yet commonsense, as I have said, understands friendship as generating obligations, a category of reasons typically categorized as moral. So one way that we can carve up the territory of reasons might be in terms of moral vs. non-moral reasons, where the former are understood as *demands* whereas the latter are not. It is certainly true that when we talk about morality, we talk in terms of duties, requirements, demands, obligations, etc. 'Ought' language

seems peculiarly at home in the realm of morality. Then we can say that moral demands fall into two categories, the impartial universalistic and the partial non-universalistic, with friendship being of the latter sort.

The problem with taking this approach to understanding morality is that it doesn't seem to make much sense to divide reasons into those that constitute demands and those that do not. A reason to do P is a consideration in favor of doing P. One ought to perform an action P, it seems, when the balance of reasons is in favor of doing P. What could it be for a reason to constitute something other than a demand? If a fact constitutes or grounds a reason, then it must have a strength relative to other reasons. Rationality, then, is a matter of acting on one's strongest reasons in the circumstances. So if there are reasons of self-interest, they must be demands in the same sense that reasons of friendship or utilitarian reasons are demands.[11] For similar reasons, I think that it is not plausible to suppose that there is more than one type of reason and yet to claim that one type, perhaps the moral (whatever we decide that is) is always overriding. Again, if reasons are considerations in favor of action, then why suppose that one type of consideration will always outweigh considerations of any other type? Consider a case of a reason of self-interest competing against a utilitarian reason: suppose that if I do what I want (or what will give me pleasure, or what will satisfy my informed-desires) that I will produce ever, ever so slightly less intrinsic value than if I made a sacrifice (a rather large sacrifice, let us suppose). Is it really plausible to suppose that if I have a reason of self-interest to do what I want that it is necessarily overridden by the impartial utilitarian reason? Is it plausible to suppose that I ought not to pull my friend into the lifeboat if a teeny tiny bit more value could be produced by doing otherwise? Is it plausible to suppose that I am always required to make sacrifices for my friend at the cost of my own good? It seems implausible to suppose that we could know *a priori* which type of reason will outweigh others.

We need to start over, I think, if we are to understand reasons and their relationship to the deliberation and actions of agents in the real world.

REASONS AND AGENTS

Discussions of reasons and deliberations have often been couched in terms of different perspectives. Thomas Nagel, for example,

claimed that the 'view from nowhere' was the vantage point from which one could recognize one's agent-neutral reasons, i.e. those reasons, the statement of which, involves no essential reference to the agent. From this perspective we are capable of acknowledging 'impersonal value.'[12]

The view from nowhere, of course, is Nagel's version of the impartial spectator, the hypothetical individual famously appealed to by Hume and Adam Smith and by the classical utilitarians as a heuristic device. This perspective is supposed to stand in contrast with the egocentric perspective of a particular agent, a perspective from which some things have a personal value to the agent that is out of proportion to the impersonal value that would be recognized from the view from nowhere. Nagel's example is the climbing of Mount Kilimanjaro: in so far as an agent very much wants to climb the mountain, her climbing of the mountain has great personal value for her. But if the agent takes up the viewpoint of an impartial spectator, she will be able to acknowledge that her climbing of the mountain has less in the way of impersonal than of personal value, and certainly no more impersonal value than anyone else's climbing of the mountain.[13]

Once we assume the view from nowhere, supposedly we will see that our own friendships are no more nor less valuable than anyone else's, in general. Of course, we might actually come to see that some of our own friendships are less worth promoting than some friendships to which we are not party. When we reassume our own egocentric perspective, our own friendships loom large, eclipsing those of other persons. In fact, the comparative value of these friendships rarely enters our minds: all that we typically consider is that *these* are *my* friendships, *those* are *not*. And now we are at the point where Sidgwick found himself at the end of *The Methods of Ethics*: practical reason seems split between the impartial view from nowhere and the partial egocentric view.

But this is a misguided way of thinking about practical reason. There is really only an egocentric perspective from which one can acknowledge various grounds of reasons that one has. After all, consider the principle of utility, our paradigm of impartiality: right action is that action which, out of all alternatives available *to the agent*, produces the greatest net sum of intrinsic value for all persons affected in the long run. Even after I assume the vantage point of the impartial spectator, I have to locate a particular agent's causal/epis-

temic position in order to determine the right action for that agent. Just seeing what impersonal intrinsic value is out there waiting to be promoted is not sufficient for understanding the rational demands upon a particular agent. So, in essence, reasons are always reason for an agent. It is, at best, misleading to talk about reasons *simpliciter*. Certain states of affairs have (impersonal) intrinsic value, but their having that value is not sufficient to draw any conclusions about anyone's actual reasons: all that we can conclude is that *if* an agent S is in a position to promote the value, *then* S has a reason to do so.

It will help to consider our ontological options here. One option would be to hold that the property of being a reason supervenes on a complicated state of affairs, the state of affairs of an intrinsically valuable fact's standing in a certain relationship to an agent, call it the 'capable of being promoted by' relation. According to this ontological picture, the grounds of consequentialist reasons always involve an intrinsically valuable state of affairs, a particular agent, and the agent's relationship to the state of affairs.

Another option is not for the ontologically faint of heart – it is not for those who always have Occam's Razor at the ready. According to this second possible ontological picture, we would hold that there is no property of being a reason *simpliciter*. Rather, there are many properties of being a reason *for S*, where states of affairs that are intrinsically valuable come to have such properties when they stand in the 'capable of being promoted relation' to S. According to this second account, reasons for agents are a function of the nature of the state of affairs not only considered in itself but also considered in its relations to rational beings in the world.[14] What is crucial is that according to either of these two ontological accounts, the very existence of reasons is a function of agents and their relations to various features of the world.

Deliberative perspectives, then, are, of necessity, always egocentric. The agent's causal position and knowledge of the world are factors that are crucial to determining the nature of her reasons for action. So now we have no split between deliberative perspectives on the world, we only have various egocentric perspectives on the world, and, for each agent, *her* reasons are a matter of *her* perspective. This does not move us to a Hobbesian or Humean conception of practical reason, as long as we acknowledge that there is objective intrinsic value in the world. But what is relevant for me is the objective value that I am in a position to promote.

Once we find ourselves securely situated in the egocentric perspective as the appropriate perspective for deliberation, allowing relationships between the agent and a state of affairs to play a role in the generation of reasons, we can see that relationships other than those of a causal/epistemic nature can also do that work. The fact that a state of affairs involves my friend's good or involves the maintenance of my friendship can have the property of being a reason for me to promote it while not having the property of being a reason for some other agent to promote it. Or, if my friend's good or our friendship is objectively intrinsically valuable, and that other agent is in a position to promote it, then she has a reason to promote it, but I have more reason to do so than she does.

Now it is true that a disanalogy remains between reasons of intimacy and consequentialist reasons. The former are, in an important sense, doubly relative to a particular agent. Recall that our consequentialist reasons are a function of our causal/epistemic relation to an intrinsically valuable state of affairs, where the property of being intrinsically valuable is a function of the intrinsic nature of the state of affairs – it in no way depends on the agent. Our reasons of intimacy are also a function of our causal/epistemic relation to a state of affairs. What kind of state of affairs? One that involves caring for a friend or sustaining the relationship with that friend. So my reasons of intimacy are a function both of (i) my causal/epistemic relation to a state of affairs, and (ii) the nature of a certain state of affairs. However, (ii) the nature of the relevant state of affairs that grounds reasons of intimacy is not intrinsic to that state of affairs, but is a matter of its relation to the agent – the relevant state of affairs will be one that involves the good of *my* friend or the sustenance of *my* friendship with a particular person.

Does this extra element of agent-relativity pose a problem? I don't see why it should. In fact, I think that when we are ensconced in our egocentric deliberative perspectives, it is self-evident that our reasons are a function of various kinds of relationships that bind us to the world: some of those relationships are merely a matter of our coming to recognize our ability to promote objective value, while others are a matter of our standing in special relationships to other persons *and* our recognizing that our ability to care for such persons and promote those special relationships. Once the complex agent-relativity of reasons and deliberation is granted, there is no reason to deny additional levels of agent-relativity, especially given that doing

so allows us to capture so much of our commonsense thinking about our deliberative and rational places in the world.

At this point we have not yet made the process of practical deliberation easy, but that should not be a goal of any theory of reasons for action. We all know that practical deliberation is *not* easy. And we have jettisoned the baggage of trying to divide reasons into overriding moral reasons and non-moral reasons. So each deliberative situation is a matter of trying to figure out all of one's reasons and then trying to determine which is the strongest in the circumstances. I am inclined to agree with Ross that there is no algorithm for doing this, that each practical deliberation is a risky situation and that one may never know for sure whether one acted rationally.[15]

So let's return for a moment to the lifeboat, where I have to decide whether to pull my friend to safety or to help a stranger. Such cases provide no counterexample to the claim that the principle of utility is the statement of one source of our reasons. Rather, what they show is that there are other sources of reasons, and that our example is one in which one of those other reasons clearly outweighs any reason to maximize objective intrinsic value. Bernard Williams famously used a similar case to attempt to undermine utilitarianism, claiming that the utilitarian would bring into deliberation "one thought too many," that the thought 'that's my friend' (or, in Williams' case, 'that's my wife') should settle the matter.[16] In fact, I am arguing, utilitarian deliberation involves at least one thought too few, and, in fact, probably several thoughts too few. The importance of objective value in no way leads to a denigration of friendship, but neither does the importance of friendship lead to a denigration of intrinsic value. If we really were occupying the view from nowhere, we could not act because we would not have a position from which to act. Given that we occupy a particular position, we acquire many diverse reasons for action. There is really nothing at all problematic about this.

One traditional worry about rendering moral reasons less than overriding is that doing so seems to leave open the possibility that it might be rational to be immoral, and, perhaps, an agent could find herself in a situation in which morality never presents itself as rational. I grant that it sounds odd to say that it is rational to be immoral, but that is because we tend to read the moral 'ought' as the all-things-considered 'ought,' and it is certainly odd to say that it is rational not to do what, all-things-considered, one ought to do – in fact, it is just plain contradictory. But it is certainly not contradictory to say that it is rational not to maximize objective value. The term 'moral' obfuscates more than it helps, and once we jettison it, we can focus on reasons, trying to determine, *a posteriori*, which reason is the strongest in any given situation.

One final point about persons for whom it might never be rational to promote objective intrinsic value: no matter what it is rational for such persons to do, it will probably remain rational for us to incarcerate or even eliminate such persons. Consider a psychopath such as Ted Bundy. Bundy told his interviewers that 'he' never imagined that anyone would miss 'his'[17] victims; after all, there are so many people in the world, and people go missing all of the time. Bundy's use of moral and sympathetic language reveals that he wasn't in complete possession of the relevant concepts; for example, whenever he talks of 'his' remorse, it is always tied to 'his' fears of being caught, never to the pain to 'his' victims and their friends and families that 'he' has caused. I have no doubt that Bundy completely lacked the capacity to recognize objective value in the world. He also lacked the ability to be genuinely intimate with another human being, as is clear from his lack of ability to understand how loved ones might miss a particular person regardless of how many people exist and how many go missing. So Bundy had no consequentialist reasons and no reasons of intimacy: he was an agent whose rational sphere was severely truncated. This had two consequences: he lacked the ability to lead a life rich in value and meaning, and we had reasons to contain Bundy, both in the service of intrinsic value and in the service of pure self-interest and love for others in the world. A perfectly rational being, such as Bundy would have been had he not sometimes acted so incautiously, might still be a sad and pathetic figure that we have every reason to both loathe and pity.

VIRTUE THEORY AND THE EGOCENTRIC PERSPECTIVE

As intimate relationships such as friendship have garnered more attention in moral philosophy, virtue theory has simultaneously experienced a significant resurgence of interest and popularity. Given what I have said so far, it seems clear why these two trends seem to go hand in hand. Virtue theory begins from an egocentric perspective: the focus, for each agent, is on the cultivation of *her own* virtuous states of character. So just as the agent has reason to be a benevolent person, she has reason to be a good, caring friend. If reasons are fundamentally reasons to have certain character traits, then virtue

theory seems to be a moral theory that avoids some of the difficulties that I discussed before.

However, I think that either (i) virtue theory cannot be understood as a distinct type of theory of reasons for action, or (ii) it is a highly implausible theory about our primary reasons for action. Consider the types of settled motivations that characterize being a good friend: aiming at the friend's good, expressing sympathy for a friend's plight, promoting the maintenance of the friendship, etc. None of these motivations have as their object the agent's own character; rather, each is aimed at the performance of some action or the sustaining or promoting of some external state of affairs. In so far as I have reason to cultivate being a good friend, that is just a matter of my having reason to be responsive to the basic reasons arising out of friendship. So there must be reasons that precede any reasons pertaining to the cultivation of virtue; in fact, the cultivation of virtue would seem to amount to nothing more than being rational. Thus, we need a theory of reasons, before we can have a theory of virtue and good character.

It might be said that, in fact, our primary reason is to develop the virtues in ourselves, and that we can somehow characterize the virtues without appealing to other reasons we have. Even if this were possible, I think that virtue theory would have become egocentric in the wrong way. All of our reasons are a function of our individual perspectives, but it seems clear to me that some of our reasons take as their object states of affairs that do not essentially involve us. For example, we have reason to promote the welfare of others regardless of what impact doing so has on our own characters. It is important to develop a theory of reasons that is egocentric in the right way: all reasons are reasons for the agent, but not all reasons are about, in some way, the agent herself. I am the agent, so my reasons are a function of my particular place in the world, but I occupy a place in a world that is not all about me, and my reasons must reflect that fact.

CONCLUSIONS

I am recommending, then, that we be careful when we use the term 'moral.' The term has a lot of baggage, baggage that can lead our attention away from the subtleties and complexities of rational deliberation. All deliberation is egocentric, but that does not mean that all deliberation is *egoistic*: all of my reasons are reasons *for me*, but not all of my reasons are *about me*. (Unless I am a psychopath, lacking the ability to recognize various normative features of the world and unable to develop meaningful, intimate relationships with other persons.) Attempts to neatly divide our reasons into moral and non-moral can mislead us into thinking that there are different deliberative perspectives. What we need to recognize is that the world contains many and diverse features of rational import, and we need to be able to simultaneously remain ensconced in our egocentric perspectives while still seeing beyond ourselves.

NOTES

1   I will use 'friend' as a general term for all persons with whom we are intimate. Some philosophers limit the use of the term 'friend' to refer to the parties to certain types of intimate relationships, for example, some do not think that the term 'friend' can appropriately be used to refer to parties to relationships that involve some sort of disparity in power and/or authority, such as the parent/young child relationship. I do not intend to be taking any stand on such issues - I am simply adopting 'friend' as a handy term to use to talk about intimates.

2   I will use 'utilitarianism' and 'consequentialism' interchangeably. I understand utilitarianism as any consequentialist theory that employs a welfarist conception of intrinsic value, i.e. some conception of intrinsic value according to which all and only the welfare of sentient creatures has intrinsic value. Utilitarians, then, may differ about the proper understanding of the concept of welfare: is it pleasure? Desire satisfaction? Informed-desire satisfaction?

3   My statement of the principle of utility is a statement of that principle in its actual consequence act- utilitarian version. Nothing in my discussion would be affected by substitution of, for example, a possible- or probable-consequence version of the principle. However, we do need to hold constant that the principle involves an appeal to a non-relativistic conception of intrinsic value, as I make clear in what follows.

4   I am, simply for expository purposes, ignoring the distinction between utilitarianism as a theory of right action and utilitarianism as a decision procedure. Of course, it might be the case that a good utilitarian agent will, in her actual thought, give greater weight to loved ones, if doing so will, in the long run, produce greater intrinsic value. Perhaps each of us ought to make ourselves into the kind of person who thinks along commonsense lines. But then, of course, our general personality traits and dispositions that give extra weight to our loved ones are justified by a line of reasoning that itself gives no greater weight to the interests of our loved ones. See Railton, Peter, "Alienation, Consequentialism, and the Demands of Morality", *Philosophy and Public Affairs* 13, 1984, pp.134-171.

5   I am, of course, also required to be impartial between myself and other persons, friends or strangers. So utilitarianism will sometimes morally require me to refuse to make a sacrifice for a friend, if maximal value would result from my not making that sacrifice. Another counterintuitive feature of utilitarianism, then, is that it will not regard all personal sacrifices for the sake of loved ones as morally commendable.

6   All else being equal, of course. Thus, for example, the stranger does not have the cure for AIDS in her head, is not someone whom I have promised to help, my friend is not terminally ill or a serial killer, etc.

7   Sidgwick, Henry, *The Methods of Ethics*, Indianapolis, IN, Hackett Publishing Company, 1981 (1907), pp.432ff.

8   For a thorough defense of a relativistic version of consequentialism, see Fumerton, Richard, *Reason and Morality: A Defense of the Egocentric Perspective*, Ithaca, NY, Cornell University Press, 1990.

9   Those philosophers who advocate a relativized version of consequentialism would disagree with this characterization of their view.

10 I am ignoring those views according to which friends are 'other selves,' and, thus, benefits to friends are benefits to the agent. I find such views highly counterintuitive. For an interesting defense of such a view, see Brink, David, "Rational Egoism, Self, and Others," in Flanagan, Owen, and Rorty, Amelie Oksenberg, eds., *Identity, Character, and Morality: Essays in Moral Psychology*, Cambridge, MA, The MIT Press, 1997, pp.339-378.

11 So I do not think that the language of 'permissions' used by many moral philosophers makes much sense: either my reason of self-interest outweighs my moral reason, and rationality demands that I act on it, or vice versa.

12 Nagel, Thomas, *The View From Nowhere*, New York, Oxford University Press, 1986, pp.164ff.

13 Ibid, p.167.

14 I have stated these metaphysical pictures in non-reductionist terms, rendering the property of being a reason or of being a reason for *S sui generis*. If one does not have my non-reductionist leanings, one could restate them in terms of reasons being constituted by certain facts rather than in terms of reasons supervening on certain facts. This option would be covered by the supervenience language if one allows that p being identical to q is a form of p's supervening on q.

15 See Ross, W. D., *The Right and the Good*, Indianapolis, IN, Hackett Publishing Company, 1930, pp.30-31.

16 Williams, Bernard, "Persons, Character, and Morality," in *Moral Luck*, Cambridge, UK, Cambridge University Press, 1981, p.18.

17 I put the relevant pronouns in scare quote because Bundy's interviewers succeeded in getting Bundy to talk about his crimes only by allowing him to describe them in the third person, as though Bundy were helping to explicate how someone could commit the horrible crimes that he, Bundy, had, supposedly, been wrongly accused of. See Aynesworth, Hugh, and

Michaud, Stephen, Ted Bundy: *Conversations with a Killer. The Death Row Interviews*, Irving, TX, Authorlink Press, 2000, p.188.