Jennie Louise   8/22/2007

Moral Demands and Not Doing the Best One Can

The problem of extreme demands is one of the most intractable in
contemporary moral theory. On the one hand, it seems that a failure to
prevent great suffering at little cost to ourselves is morally wrong; given
the amount of suffering in the world and the comparatively trivial nature
of the requisite sacrifices, this intuition demands that we give up quite a
lot. On the other hand, it *doesn't* seem to us that we act wrongly in
living lives characterised by only moderate sacrifice, in which our time
and resources are disproportionately used to benefit ourselves and those
close to us. These two intuitions are extremely difficult to reconcile
within any moral theory that recognises a duty to promote the general
good. In this paper, however, I will suggest one possible way of doing so.
My suggestion requires taking a closer look at the way in which the
demand to the promote the good is derived: specifically, at the way our
option set is characterised and the information that we take into account
in weighing these options. I will suggest that there are certain
assumptions it is plausible to make regarding the relevance of
information about our own and other agents' actions, and that once
these assumptions are made, we can see how permissions may be
derived within the framework of good-promotion.

1. The Problem of Extreme Demands

Since the problem of Extreme Demands is well known, I do not intend to
spell it out here. I will, however, briefly outline the assumptions and
framework I will be using in this paper. Firstly, I will be discussing only
so-called *demands of beneficence* – that is, demands that we contribute
to the welfare of other agents. Although some philosophers have pointed
out the need to address the potential for equally extreme demands

arising from duties and rights, this is beyond the scope of the present paper.[1] Secondly, I am assuming that the demands of beneficence take the form of a *pro tanto* reason to promote (neutral) value: specifically, the value of human welfare.[2] Given the state of the world, promoting the welfare of others requires me to make very large sacrifices in terms of my own welfare. Our commonsense moral intuitions, however, tell us that such sacrifices, while morally admirable, are not required of us. In fact it is morally permissible (that is, neither morally required nor morally prohibited) for me to favour my own interests to a very significant extent.

The problem of extreme demands is therefore that of finding room, within a neutral-value-promoting framework, for a permission to choose options that are less than the best available to us. Although a certain amount of sacrifice may be morally required in order to promote the good of others, this cannot rise beyond a relatively moderate level: for example, that which would be my fair share of a sufficient contribution when distributed amongst all agents to whom the demand applies. Beyond this, a plausible moral theory will neither require nor prohibit anything: agents will be permitted either to devote their left-over time and resources to their own welfare, or to make extra contributions to compensate for the non-compliance of others.

A third assumption, which is perhaps more controversial, is that the problem of extreme demands – and hence, any solution – will arise in the context of decisions that have fairly wide scope: that is, decisions about what kind of life I am going to try to live, or what level of regular

---

[1] I believe, however, that the account I develop here may contain the resources to give such an explanation.

[2] By neutral value, I mean value which is not relativised to agents (e.g., the value of *my* welfare) or to times (the value of *present* welfare).

contributions to the worse off I will commit myself to, rather than individual decisions such as, for example, whether to use the three dollars in my hand to buy a cup of coffee or donate it to charity. The kind of permission I am looking for, then, is a permission to choose to live a life characterised by moderate sacrifice, in which I exhibit disproportionate partiality towards my own projects, and towards those with whom I have a close personal relationship. Whether permissions in individual choice situations (what should I do with *these* three dollars?) follow from a permission to live a life of moderate sacrifice is a subject for another paper.

2. Options and Outcomes

I am therefore concerned with a choice situation about the level of demands to which I ought to commit myself. I might choose to commit myself to a life of extreme sacrifice, giving up most of my projects and any personal relationships that would require disproportionate partiality. Or I might choose to use my resources to live an ordinary life, in which I maintain personal projects and relationships, thus contributing only a modest amount of my resources to the welfare of others. From the point of view of demands to promote the good, it seems obvious which I should choose; no plausible argument can be made for thinking that the neutral value of my projects and relationships – that is, their contribution to general welfare – comes close to the value that would be realised if the resources used to sustain them were instead devoted to improving the welfare of the worst off people in the world.

One thing to note, however, is that the choices have been described in a way that assumes particular outcomes for each. In particular, it is assumed that whatever kind of life, or course of action I commit myself

to, I will follow through on that commitment. But this is not a realistic way to describe my actual choice situation. Of course, if I commit myself to a particular course of action, it is at least *possible* for me to follow through, and I make the commitment in the expectation or hope that I will do just that. But this is by no means a certainty, even if we ignore the trivial fact that my circumstances could be altered by factors beyond my control (for example, my being run over by a bus before I donate any money to charity). More importantly, whether my present choice has the specified outcome depends upon my own future actions – in later choice situations, I have to choose in a manner appropriate to my earlier decision. Thus, whatever choice I make now, there is in fact a *range* of possible outcomes depending upon what choices I make subsequent to the initial choice.

How should the fact of a range of possible outcomes inform my choice between options? The two main positions on this question are commonly referred to as Actualism and Possibilism. Possibilism holds that, of a set of options, the best (and hence the right) one is the one that has the best possible outcome. Actualism, on the other hand, determines the value of each option by reference to the outcome that *would actually* occur were that option to be chosen; hence, the best option is the one with the best actual outcome.[3] Clearly, the two will deliver different judgements in cases where the actual outcome of choosing an option will not be the best possible outcome of choosing that option.

In the case of demandingness, we can assume that the best possible outcome of my making a commitment to a life of sacrifice is that I

---

[3] These definitions are given in Frank Jackson and Robert Pargetter, "Oughts, Options and Actualism," *The Philosophical Review* vol.95 no.2 (April 1986), 233.

successfully follow through on that commitment; although my own welfare is significantly lower as a result, the amount of general welfare I produce is significantly higher. But we can suppose that, if I am like most other people, what will actually happen subsequent to my making such a commitment is that I will fail to live up to it; following a great deal of suffering and guilt on my part, and perhaps irreparable damage to my projects and relationships, I would make only a modest contribution. The best possible outcome of my making a commitment to the modest standard, however, is far closer to what would be the actual outcome: I would succeed in following through on that commitment (say, doing my fair share), while maintaining my own welfare at a higher level.

If I am choosing on the basis of the best possible outcome, then the best possible outcome of a commitment to a life of sacrifice (a significant contribution to the general welfare) is better than the best possible outcome of a commitment to doing my share (a modest contribution to the general welfare); thus, it seems that I ought to make the former commitment. But of the outcomes that will actually result from my choice of each option, that associated with doing my share (a modest contribution with higher welfare for myself) seems better than the outcome associated with the life of extreme sacrifice (a modest contribution with a significant amount of damage to my welfare); therefore, I ought to adopt the latter option.

The question, then, is whether it is legitimate to allow my choice to be influenced by my prediction that I *will* act in a certain way subsequent to choosing, or whether the only thing that should be relevant is what I *can* do subsequent to choosing. According to Possibilists, if I *can* bring about the best possible outcome – if I do not expect that I will be rendered incapable of responding to my reasons in my subsequent choice

situations – then it is not appropriate for me to base my current choice on a prediction that I will in fact fail to respond to them. To choose in this way is to fail to treat oneself *as an agent*. According to Actualists, however, to fixate upon the best possible outcome regardless of the fact that it will not in fact happen is arbitrary and irrational, especially when this also means ignoring the likelihood that this choice will in fact bring about the *worst* possible outcome. I will have something to say about this debate later on. First, however, I would like to examine the two positions in the context of the demands of beneficence.

3. Possibilism, Moral Demands, and the Agency of Others

If, in a value-promoting framework, I ought to choose the option with the best possible outcome, then promoting the good apparently requires me to commit to a very demanding standard. As an agent – that is, as someone who has the capacity to respond to reasons, and who is aware that I possess that capacity – I cannot, when deliberating about my choices, make a decision based upon an assumption that I will not actually respond to those reasons. I am therefore not permitted to choose a course of action in which I maintain my projects and attachments, if I *can* meet a more demanding standard.

The claim that extreme sacrifice has the best possible outcome is made on the basis of a number of background facts. Some of these – the fact that many people have an unacceptably low standard of living, and that large contributions from the affluent can increase their welfare by far more than they decrease the welfare of the contributors – ought to be uncontroversial; at any rate, they will not be in question here. However, another important background fact is an assumption about the non-compliance of other agents. The demand to promote general welfare is

addressed to any agent who is in a position to do so; if every such agent in fact responded appropriately, the demands on any individual agent would be far less extreme (or at least, it would be much harder to object to them). Intolerably extreme demands arise in part because the vast majority of other agents are not meeting their obligations.

But if it is illegitimate for me to allow assumptions about my future actions to enter my deliberations – and the reason for this is that such assumptions involve a denial of my own agency – then the claim might also be made that it is inappropriate to make assumptions about the actions of *other agents* in making my choices. If the fact that I am *capable* of responding to my reasons rules out my choosing on the basis that I *won't* respond to my reasons, then doesn't the fact that others are capable of responding to their reasons similarly rule out my choosing on the basis that *they* won't do so?

If, in making my calculation, I must disregard any information about what other agents will do as well as information about what I will do, then the picture changes significantly. In this case, the best possible outcome of my doing just my share is that I succeed in following through, and all other agents do as well. The best possible outcome of my adopting a life of extreme sacrifice, on the other hand, is that I succeed in following through, and that others also do at least their share. And the latter outcome seems worse in at least one respect than the former: that is, my welfare is significantly affected in the latter. Thus, adopting a life of extreme sacrifice, whatever its actual consequences, will not produce the best *possible* option. It may appear, then, that although one aspect of Possibilism pushes in the direction of extreme demands, the other pulls away from it.

Jennie Louise   8/22/2007

An argument similar to that above has been made by L.J. Cohen.[4] And as Humberstone points out, this position seems initially appealing, due to "its refusal to let those who can and should pay 'off the hook' by regarding others as obliged to act, in effect, as if the non-payers were not there at all. It seems crucial that the defaulters are in fact psychologically capable of contributing."[5] However, it is not in fact a plausible position. Humberstone's objection is that "to plump for the action which yields the best world only together with actions on the part of others which they will in fact not perform is typically to act less well than one might, given one's circumstances."[6] In other words, it seems intuitively obvious that we ought to take into account what other agents will do when deliberating about what *we* ought to do.

This intuition exists in the case of requirements to promote the good as well, as can be seen by considering a simple example. Suppose that there are two victims drowning in a lake, and there are two people in the vicinity, each aware of the victims' plight and capable of performing the rescues. The best possible outcome is clearly for each person to rescue one victim: it is fairer, there is less risk to each person from overexertion, and the victims are rescued more quickly. But it would be absurd to claim that the first rescuer ought to ignore the other's inaction and rescue only one victim, on the grounds that this is the option associated with the best possible outcome.

---

[4] L.J. Cohen, "Who is Starving Whom?" *Theoria*. The interpretation in terms of the actualism/possibilism debate is made by Lloyd Humberstone, "The Background of Circumstances," *Pacific Philosophical Quarterly* vol. 64 (1983): 19-34.

[5] Humberstone, "The Background of Circumstances," 28.

[6] Humberstone, "The Background of Circumstances," 28

The notion that we ought to ignore what others will actually do might be defended at this point by the argument that in this case, the outcomes are not merely probable but certain. If I do not rescue the second victim, they will die; there is no longer any possibility of the outcome in which the other person performs the second rescue. If the other would-be rescuer *refuses* to do their part, then I am justified in discarding that possibility from the set of possible outcomes. The best possible outcome then becomes the one in which I rescue both victims, and I therefore ought to do that.

The problem with this reply is that such a standard, if sufficient to deliver the desired result in the rescue case, will almost certainly also lead to extreme demands to promote the good generally. If the other potential rescuer's inaction *thus far* is enough to negate the possibility of their acting in the rescue case, then the inaction of others thus far in the face of global poverty is also sufficient to negate the possibility that they will suddenly *now* do their share to alleviate it. And this will lead back to extreme demands: I will have to take the inaction of others into account and commit myself to the course of action with the best possible outcome given that assumption.

In any case, however, this approach is not going to deliver permissions. Although the moral rightness of the choice to live a life of only modest sacrifice would fall out of an account which allowed assumptions about what other agents were going to do, this would be a *requirement* to live such a life, not a permission. Even more counterintuitively, making more than the modest sacrifice will be prohibited, since doing this will be incompatible with the realisation of the best possible outcome.

4. Actualism and Extreme Demands

We have seen that it is intuitively appropriate for me to consider what other agents will do in deliberating about my options. But what about considering what *I* am actually going to do? As noted, assuming that other agents will not comply with demands of beneficence leads to extreme demands for myself. However, if *I* will not follow through on a commitment to meet such demands, then perhaps this also ought to influence my choice. As Frank Jackson notes, "some actions are such that they only have good results if they are followed up in the right way . . . . In all these cases it is better not to start if you are not going to follow up in an appropriate way."[7] So if I will not follow through on any commitment to live a life of extreme sacrifice, and if it is appropriate for me to consider this in my deliberation, then I ought to adopt some choice that will have better actual results.

Unlike the case of other agents, however, it is not so obvious that my own subsequent actions should influence my choice. If these subsequent actions are under my control and if I am capable of responding to reasons, then it is just not acceptable to choose an option that is less than the best because I won't actually respond to my reasons. Lloyd Humberstone's example is of a doctor in a concentration camp performing sadistic experiments on prisoners. Ought the doctor to give his next victim an anaesthetic on the grounds that he is actually going to torture them, and it is better that the torture is less painful? It seems obvious that he ought not to administer the anaesthetic, because he also ought not to be torturing the victim. That he is going to torture the

---

[7] Frank Jackson, "Decision-Theoretic Consequentialism and the Nearest and Dearest Objection," *Ethics* vol. 101 no.3 (April 1991), 480.

victim anyway is not something he should consider in deciding what to do.[8] By the same token, we might think that a predicted failure to live up to requirements to promote the good is neither here nor there when we deliberate about whether to commit ourselves to doing so.

In response to this kind of objection, defenders of this view point out that obligations are relative to sets of options. And while the lower standard might be my obligation relative to one set of options, this does not make it my *only,* or all-things-considered, obligation.[9] Although I ought to choose the option that has the best actual outcome, this does not mean that I do not *also* have an obligation to choose the option with the best possible outcome *and then bring about that option.* If my option set is

---

[8] Humberstone, "The Background of Circumstances," 22-23.

[9] See the discussion in Jackson and Pargetter, "Oughts, Options and Actualism," 240-241. The stipulative nature of option sets also allows this account to avoid the objection that, since the set of all options available to the agent is indefinitely large, it is not possible for us to know which option is the best. The objection, put to me by Tim Chappell, requires a longer reply than is appropriate here; however, it does not apply to an account which holds that bestness/rightness holds relative to some *specified* option set. Chappell's concerns about stipulating options sets are (a) that this would mean giving up any claim that Consequentialism's moral judgements are based upon a comprehensive survey of available options, and (b) that it allows the Consequentialist to set up option sets in such a way as to suggest the desired answer. See Chappell, "Option Ranges." Regarding (a), my inclination is to bite the bullet. Regarding (b), it seems to me that, even if this were true (of which I am unconvinced), the critic could remedy the situation by proposing specific alternatives to be included in the option set. The Consequentialist account developed here already acknowledges that what is best relative to one set may not be best relative to another; the question is then to determine if there is a 'definitive' option set for the purposes of evaluating the agent's choice. For more discussion of this point, see Frank Jackson, "Decision-Theoretic Consequentialism and the Nearest and Dearest Objection," and "On the Semantics and Logic of Obligation," *Mind* vol. 94 no. 374 (April 1985): 177-195.

described as {committing myself to extreme demands; committing myself to modest demands}, then I ought to choose the latter because what will actually happen subsequent to my making that choice is better than what will actually happen subsequent to my choosing extreme demands. But if my option set consists of {committing myself to extreme demands and following through, committing myself to moderate demands and following through}, then the outcomes are already built into the choices, and I ought to choose the more demanding option. So I have an obligation to choose the less demanding standard because that will have the best actual results, but I also have an obligation to choose the more demanding standard and follow through, because this is the best outcome that I *can* bring about. As Jackson and Pargetter point out, "we must not let the mere fact that someone would not do something show that they could not or ought not."[10]

But this then seems to indicate that not only do we not have permissions on this account, we do not even have moderate demands. I have not shown that, on the Actualist account, only the moderate demand exists: rather, a moderate demand has been *added* to the extreme demand, which is still there, and still demanding.[11] If it is what will actually bring about the best consequences, then I am *required*, rather than permitted, to devote myself disproportionately to my projects and attachments. And if I instead manage to meet a higher standard, this is not praiseworthy supererogation: in *that* case, the lesser obligation drops out, because it is not then true that I won't actually meet the higher standard, and the higher standard is precisely what is required of me.

---

[10] Jackson and Pargetter, "Oughts, Options and Actualism," 240.

[11] Thanks to Garrett Cullity for pointing this out.

5. Predictions, Probability and Agency

I think that appearances here may be somewhat misleading, however. We need to look more closely at what it means to claim that a particular outcome will actually result if a certain option is chosen. One thing this had better not mean is that, it is *true now* that, for example, I will ϕ later if I choose option X now: if this is the case, then it is hard to see how I can assert that it is nevertheless *possible* for me to ψ subsequent to choosing X. Rather, the best interpretation of the claim that if I choose X, I will later ϕ, is as a *prediction* of the most likely outcome. This prediction is made on the basis of self-knowledge, facts about the situation and perhaps inductive reasoning from my past behaviour.

Thus, we might think that my actions subsequent to any choice should inform my decision via a sort of Expected Utility model. The claim that I will ϕ if I choose option X is simply an assertion that, given the available facts, my ϕ-ing is highly probable (although it cannot be a certainty). But this is certainly not a standard version of Expected Utility theory: the claim is not that I ought to calculate the probabilities of all possible outcomes of my choosing an option, then multiply by the utility each outcome would produce. Rather, the claim is that the utility of the option *is* the utility of the highly likely outcome.

But of course, not all options have one outcome that is highly likely. And this may explain the choice of examples by each side in the debate concerning the legitimacy of taking into account your own future actions. Actualists tend to use examples in which there is a particular outcome, consisting of the agent's acting in a certain way, which is *much* more likely than any other outcome, given facts about the agent's past behaviour and their character and dispositions. Possibilists, however,

tend to use examples in which it is not as plausible to assert that there is a very high probability of the agent's subsequently acting in a particular way: in which there does not seem to be, for example, weakness of will, strong feelings or attachments, or anything else that would make it difficult for the agent to act for the best.

The point of the objections against Actualism, however, is not to show that there are *some* cases in which the relative probabilities fail to support a claim about what the agent will do. The objection is to the idea of using probabilities of any kind in deliberating. The claim is, in other words, that it is *inappropriate* to appeal to probabilities as a guide to action, no matter how high or low.[12] Indeed, we might go further and say that it is inappropriate even to *assign* probabilities to one's own free future actions. To do this is simply to ignore or deny the fact that one is an *agent,* and that the role of deliberation is to *decide* what one will do. No matter how unlikely any given outcome, if I *can* bring it about, a decision to do so might render these probabilities irrelevant. I cannot, in other words, treat my own actions as merely another empirical phenomenon to be predicted, because *agency confounds the assignment of probabilities* to outcomes which are constituted (even partly) by one's own actions.

In fact, there are two possible ways to make this claim. One is that, although probabilities can be assigned to my possible future actions, it is inappropriate to pay attention to them in my decision-making. A second, stronger, form of the objection is that the fact of my agency renders probability assignments *impossible*, at least from my own perspective.

---

[12] The exception is of course where that probability is *1* (or 0); this is acknowledged to be relevant to deliberation by Possibilists as well as Actualists.

Either of these claims will support the argument that I cannot use a prediction about my own future actions to justify my choice of a less good course of action. According to the weaker claim, although it is highly probable that I will act in a particular way subsequent to a choice, I should not consider this fact in deciding which choice is best. According to the stronger, claim, it is not even true that my acting a particular way has a determinate probability.

I think that there is something to these objections. It does seem incompatible with one's concept of oneself as an agent to adopt a predictive stance towards oneself. As Richard Moran writes,

> what is wrong with the 'direction of gaze' here, the shift to the theoretical or empirical perspective, is that it suggests that [one's] reflection on [one's] best reasons for belief or action still leave it an open question what [one] will actually end up believing or doing. This is not a stable position one can occupy and still conceive of oneself as a practical and theoretical deliberator. One must see one's deliberation as the *expression and development* of one's belief and will, not as an activity one pursues in the *hope* that it will have some influence on one's eventual belief and will."[13]

On the other hand, there are some circumstances in which we do *not* think that an agent ought to ignore a high probability that they will act a certain way; nor is it plausible to deny that their acting this way is highly probable. If my past decisions to ϕ have always met with failure, and there is nothing different about my current situation, then it doesn't seem that I should treat my present decision to ϕ as if it will succeed unproblematically. Moran illustrates this with the example of Sartre's gambler, who resolves to stop gambling but who also knows that he will most likely fail to do so. As Moran writes, "A certain 'realism' demands

---

[13] Richard Moran, *Authority and Estrangement: An Essay on Self-Knowledge* (Princeton University Press, 2001), 94.

that he acknowledge that, in his current condition, his 'decision' to quit does not fully answer the question of what he is actually going to do . . . . He must see his current decision and his empirical history as each delivering incompatible answers to the question of what he is going to do . . . . Yet, he must also see that he cannot simply dismiss the claims of either perspective."[14]

Regarding the weak claim, this seems to mean that some attention to probabilities is warranted in cases where there is strong evidence in favour of a particular outcome. Regarding the strong claim, it seems that *some* assignment of probability is possible, albeit only of a very coarse-grained kind. Where the rough probability of a particular outcome is *obvious* (very likely, highly unlikely), it may be appropriate for an agent to allow this fact to influence their deliberations. And yet, predictions about one's future actions can never be substituted entirely for deliberation on the basis of what one is capable of doing as an agent.

6. Probability and Permissions

So what does this have to do with extreme demands? What we are looking for is an account of requirements to promote the good that permit, but do not require, sacrifices beyond a given level, and that also permit, but do not require, a life in which I devote my resources disproportionately to my own projects and relationships. I have claimed that, in deliberating about demands to promote the good, it is not legitimate to ignore my predictions about what other agents will do. As far as my own actions are concerned, although I cannot take a stance

---

[14] Moran, 162.

that is *merely* predictive, I also cannot ignore predictive information when it is very strong and well-grounded.

In the case of demands to promote the good, I can choose to adopt a course of action involving extreme sacrifice, or to adopt a course of action involving more moderate sacrifice. Given everything I know about myself and my past behaviour, as well as my knowledge of human nature in general, it seems that I ought to take into account in my deliberations the overwhelming likelihood that I won't manage to follow through on a more demanding commitment. However, it would not do to suppose that I should attempt to calculate the exact probabilities of my carrying out various courses of action subsequent to my choice. As we have seen, it is disputable whether the calculation of such probabilities is even possible, and even if it is, this is not a perspective that I can take while still regarding myself as an agent deliberating about what to do.

How, then, ought we to take predictions about our own actions into account? My suggestion is that, in considering the value of my options, I ought to discount the value of each option in a way that is commensurate with my rough estimate of the probability that the best outcome will occur. But the perspective of agency probably precludes anything other than a very crude estimate: in general, the stronger the evidence that I will not succeed, the more I should discount the value of that option, but I will not be able to say precisely how much. So from the agent's own point of view, the value of their options will not be able to be precisely determined. And if this is correct, we may after all be able to account for permissions regarding promoting the good.

It has in fact long been recognised that permissions can be generated within neutral-value-promoting frameworks: a permission is generated

whenever an option set contains two options that are equally good, and (jointly) better than all of the other options in the set. If I cannot choose both, I am permitted to choose either. This kind of permission, however, has usually been regarded as trivial and not of any practical importance, since such cases will arise only very rarely, and certainly will not occur in relation to beneficence: almost all parties to the debate agree that the neutral value of my living a life of great sacrifice is much greater than the neutral value of my living a life of moderate sacrifice.

Even when we take probabilities into account, it does not seem that a permission will be generated. It may bring the values of the two options closer to one another: the overall value of choosing the life of moderate sacrifice (moderate value, with a high probability of its occurrence) may not be vastly different to the overall value of choosing the life of great sacrifice (very high value, with a very low probability of its occurrence). But unless they are *exactly* the same, permissions will not be generated; rather, what will result is still a requirement to choose a particular option.

This, however, is where the inability to precisely specify probabilities or expected utilities becomes useful. What I have claimed in the preceding discussion is that requirements to promote the good are the kind of requirements for which I ought to take into account the likelihood of my failure to meet extreme demands. But because I cannot assign anything more than very rough probabilities to the relevant outcomes, the adjusted value of my options will likewise only be able to be specified in a very rough way. And this means that, in a wide range of cases, I will not be able to make a judgement about whether one option is better than the other. And if I cannot differentiate between the value of the options, I do

not have a basis for choosing one rather than the other. What I would like to suggest is that this will give me a permission to choose either.

My account therefore implies that permissions regarding promotion of the good arise from the interaction of two factors: firstly, the robustness of the empirical facts regarding what I will in fact do subsequent to choosing a course of action, which makes it appropriate for me to consider these facts in deliberation. The second factor is my inability to take these facts into account in anything other than a very vague way, so that the value assigned to my options is approximate. This is consistent with both the weak and strong claims above: according to the weak claim, I cannot put myself in a position to determine probabilities without undermining my status as a deliberating agent. According to the strong claim, I cannot assign more than approximate value because it is not possible to determine probabilities regarding my own actions. Either way, my inability to determine which option is better than the other, when the value of each is adjusted by the likelihood of my realising it, results in a permission to do either.

There are a few nice features of this account. Firstly, it does not rely upon any form of the appeal to cost, or any claim that agents are allowed to assign a disproportionate weight to the value of their own projects and attachments in their decision making. Secondly, it gives a kind of explanation for why we think that adopting a life of extreme sacrifice is *praiseworthy.* The successful choice of such a life shows that the agent has been able to transcend the empirical facts about their character and past patterns of behaviour; it involves taking oneself seriously as an agent, and the exercise of considerable willpower and self-control. But the praiseworthiness of this course of action does not show that all agents should ignore the empirical evidence and embark upon

demanding courses of action based only upon the possibilities presented by their agency.

7. Objections and Problems

The argument I have presented above is fairly broad and lacking in detail; and a number of problems still need to be sorted out before deciding whether the theory is viable. In this section I will mention some potential objections. Some of these are fairly easily answered; others are more difficult, and more work will be required to determine their significance.

One argument which might be made against this theory is that, as an elaboration upon the Actualist account, it suffers from pretty much the same problem as Actualism: namely, that it imposes a moral standard that is far too lenient. For example, why assume that any given agent is highly likely to follow through on a commitment to a life of *moderate* sacrifice? Would this account not generate a permission for me to do *nothing at all* to help others, on the grounds that I am so lazy and selfish that the probability of my doing anything more is very low? On the other hand, much more could be demanded of someone who is not captive to their dispositions: if there is not a strong probability that I will fail, the more demanding option will obviously be better, and I will be required to choose it.

I think that my account will lead to stronger demands being placed on those whose decisions are generally more efficacious in producing the decided-upon outcomes. It also does not seem that there will be a bright line between cases in which it will be appropriate to appeal to predictive facts, and cases in which it will not. Rather, the boundary of

permissions will be rather vague. But there is probably a level of ability to render one's decisions effective that is a minimum standard for considering someone an agent at all. Anyone who is genuinely an agent cannot plausibly assert that a decision to do even a minimal amount to help others would not be efficacious.

Another line of argument here would be to claim that the legitimacy of making assumptions about one's future actions is dependent on the existence of very strong empirical evidence, not only about oneself, but also about normal human agents. That is, if a normal agent would have no problem in following through on a particular commitment, it will not be legitimate for any individual agent to deliberate on the basis of a prediction that they *will* have a problem following through on that commitment. This would block appeals to my laziness and selfishness to generate too-generous permissions, while also retaining the moderate-sacrifice permission, for it is obviously very difficult for normal agents to successfully commit to a life of extreme sacrifice. This line of argument would require much more work, since the question of what normal agents would or would not succeed in doing is likely to be a difficult and complex issue. I will not offer any general principles or solutions here. However, I will submit that if there is *any* case in which the evidence for predictions about normal human agents is strong, it is that of promoting the welfare of distant strangers.

An objector might point out that there is another aspect to the problem they have presented: there may be some cases in which the empirical evidence about one's future actions is so strong, and so much rides on the outcome, that it may be clear that taking the second-best option is better, thus resulting in a requirement to do this rather than a permission. In such cases it would be wrong to choose the perspective of

agency. For example, suppose I am considering making a commitment to donate my life savings to a particular charity in order to save people from some natural disaster. As a result of my commitment, the charity would consider its fundraising objectives met, and would stop soliciting donations, so that if I did not in fact follow through on my commitment, many people would die due to the delay in calling for more money. If I have extremely good reason to believe that I will not in fact live up to my commitment, then it would, on this account, be wrong for me to make it. But I submit that this is in fact a fairly plausible claim. The more that is at stake in any decision, the less defensible it is to take risks.

I believe, then, that the first objection will not prove fatal to my position. However, a more important issue for my account is the question of *objective* permissions. The account I have given is one of subjective permissions, since it relies on the impossibility or inappropriateness of making predictions about one's own behaviour. But the same problems do not exist regarding predictions about the behaviour of *other* agents: as we have seen, we should take these into account in our own deliberations. Indeed, Jackson and Pargetter try to support their claims about Actualism by asking us to consider what advice a third party would give in cases where the probability of my acting a certain way is high.[15] But if a third party can take a purely predictive stance, then probability assignments are also appropriate from the third-person perspective. And even though they may not be very precise, they will be more precise than those that can be given by the agent themselves. And this means that there will be a large range of cases in which a comparison between the value of options is possible from a third-person perspective even though it is not possible or appropriate from a first-

---

[15] Jackson and Pargetter, "Oughts, Options and Actualism," 237.

person perspective. In these cases there will be subjective permissions but not objective permissions.

I am not entirely sure what to say to this objection, or indeed how much of a threat it is to my account. But I think that, when we are talking about demands to promote the good, the first-person stance is most relevant for the simple reason that I am deliberating about what I ought to do. A third-person perspective, in which the probability of my following through is more easily determined, will often result in an objective requirement for me to adopt the second-best option. But from my own perspective as an agent, these probabilities, and hence the requirement resulting from them, does not seem relevant to my choice. It does seem clear, however, that any further development of this account will have to incorporate a more detailed solution to this problem.

8. Conclusion

Accounts of demands to promote the good have tended to overlook the role that probabilities play in determining our obligations. Indeed, to some extent, the opposite extreme has been the case: our ability as agents to bring about any possible outcome through decision has obscured the need to concentrate upon our limitations. It is a tricky and delicate question exactly how seriously we ought to take these limitations: if we take them too seriously, we are not trying hard enough and are denying our agency, but if we don't take them seriously enough, we are just being irresponsible.

What I have tried to argue for here is the claim that, where demands to promote the good are concerned, the perspective of agency and the

perspective of prediction are equally legitimate, given the difficulties of assigning precise probabilities to the outcomes involved. It is true that we *can* do better if we try; on the other hand, everything we know about ourselves and about the behaviour of human agents in general suggests that this is highly unlikely. Given these facts, and the much higher probability of succeeding in doing less, it would seem equally good – roughly – not to do the best we can.