CANADIAN JOURNAL OF PHILOSOPHY
Volume 37, Number 1, March 2007, pp. 153-172

Brute Requirements

Critical Notice

JOSHUA GERT, *Brute Rationality*. New York: Cambridge University Press 2004. PP????[1]

Brute Rationality is a remarkable book. It is elegant and engaging, and it aims to deliver straightforward solutions to problems not addressed by other theories of practical rationality. The theory of normative reasons developed in *Brute Rationality* is both detailed and novel. According to Gert, most of the literature on this topic has gone badly wrong by assuming that reasons can only have what Gert calls a 'requiring role.' To assume that reasons have only a requiring role is to be committed to the following type of inference: if the fact that I have a beautiful voice is a reason for me to sing now, I am rationally required to sing now, unless there's a stronger reason for me to do something that is incompatible with my singing now. According to Gert, theories of practical reason committed to this assumption have overlooked the fact that reasons can also have a justifying role. It might be the case, for instance, that my having a beautiful voice justifies, but does not require, my singing on certain occasions. If so, Gert maintains, the following will be true. First, this reason is capable of making a certain action (say, my singing now)

1 All page references in the main text of the article are to this book. An earlier version of this paper was read at an 'Author Meets Critics' session at the 2006 Pacific Division Meeting of the APA. I would like to thank the audience for their questions and comments, and especially Joshua Gert for his insightful replies on that occasion. I would also like to thank Tom Hurka, Jonathan Peterson, Fred Schueler, and two anonymous referees for this journal for comments on earlier drafts. Research on this paper was partly funded by a grant from the Social Sciences and Humanities Research Council of Canada.

rational where it would have been irrational in the absence of this reason.[2] But in contrast with most views of rationality, my failing to act in accordance with this reason would not render my actions irrational *even in the absence of countervailing reasons*. That is, since my reason to sing does not have a requiring role, if I don't sing, I am not failing to do anything that I am rationally required to do. So no matter what other reasons I may or may not have, I am not being irrational when I choose to let my beautiful voice remain dormant. Moreover, reasons can have differing strengths along the requiring and justifying dimensions. Roughly speaking, my reason to see the *aurora borealis* can be very weak along the requiring dimension, since I might not be rationally required to do almost anything to see them, but very strong in the justifying dimension, since it might be the case that I would not be irrational if I were to use up all of my vacation time and a good chunk of my savings to take a trip to the Yukon in order to witness this phenomenon. Although this might seem at first a minor correction to orthodoxy, Gert makes a compelling a case that wide-ranging consequences follow from accepting that reasons can have these two distinctive roles, and in the course of making the case for his view, Gert touches on almost every major debate in the area. For instance, Gert argues for (i) a new way of making a distinction between objective and subjective rationality, defending a reliabilist view of the latter (ch. 7); (ii) against the seemingly obvious claim that practical reasons necessarily motivate an agent insofar as she is rational (ch. 3); (iii) for the claim that it is neither rationally required nor irrational to act morally (ch. 2); (iv) for a functional analysis of reasons in terms of their systematic contribution to the rational status of actions (ch. 4); and (v) for a substantive theory of objective reasons in terms of harms and benefits to oneself and others, in which typically harms and benefits to oneself are strong reasons along the requiring dimension but weak reasons along the justifying dimension, and harms and benefits to others are weak reasons along the requiring dimension but strong reasons along the justifying dimension (ch. 5).

Gert makes original contributions to all the topics he discusses. However, my focus here will be on what is perhaps the central idea of the book: the claim that philosophers have overlooked the need to distinguish between these two different roles that reasons can play, a justifying and a requiring role.[3] According to Gert, as long as we fail to appreciate

---

2 I am assuming, of course, that there are no other reasons that would make rational for me to sing now.

3 See *Brute Rationality*, 76 and many other passages.

this distinction, we will be incapable of accounting for seemingly plausible judgments about the rationality or irrationality of actions. Gert argues that this very basic and simple distinction allows us to explain 'phenomena' of rationality that no other theory can explain. His arguments are extremely compelling. However, I want to suggest various ways we can try to account for the phenomena, and I'll argue that these strategies have important advantages over Gert's account in terms of the distinction between requiring and justifying roles.

I

According to Gert, the same reason can have widely divergent requiring and justifying strengths. Here is his conception of what makes it the case that a consideration plays a justifying role:

> A consideration is a reason if it can make it rationally permissible to perform actions that would be irrational without it. (66)

And as for the requiring role we get the following definition:

> A consideration is a reason if it can make it irrational to do something that would, without that consideration, be rationally permissible. (67)

As Gert makes clear, it is not the case that 'some reasons play a merely justifying role because they are comparatively weak.' Reasons are assessed for strength in their different dimensions. So the comparative strength of reasons R1 and R2 along the justifying dimension is determined as follows:

> (i) R1 would make it rationally permissible to do anything that R2 would make rationally permissible to do.
>
> (ii) R1 would make it rationally permissible to do some things that R2 would not make rationally permissible to do. (68)

A similar definition, *mutatis mutandis*, determines the comparative strength of a reason along the requiring dimension.

As we said above, the same reason can differ widely in terms of its strength along the requiring and justifying dimensions. Certain altruistic reasons, for instance, are likely far down on the spectrum of requiring strength, but far up on the spectrum of justifying strength. It would not be irrational for me simply to ignore that the $200 I used to buy my niece a new iPod could also be used to save 10 children from starvation (the potential prevention of the death of 10 children requires very little from

me), but it would also not be irrational for me to risk life and limb to save 10 children from starvation (the potential prevention of the death of 10 children can justify quite a lot). It is important to note that although reasons can differ in strength along the justifying dimension, the theory does not deem an action irrational in virtue of the fact that there was a stronger justifying reason to do something else. To think that one must always act on the stronger justifying reason is to confuse justifying and requiring roles.[4] As long as the reason is sufficiently strong along the justifying dimension to make a certain action rational, the fact that there are other reasons that are even stronger does nothing to change the rational status of the action. The fact that an action will save the lives of ten children has much more justifying strength than the fact that it will provide me a modicum of aesthetic pleasure. But since the latter reason suffices to justify my spending $200 in order to buy my niece an iPod, the fact that the same $200 could be used to save ten children does not affect the rational status of my buying my niece the iPod. Again, this does not imply that buying my niece an iPod has the same justifying strength as saving 10 children. Saving the lives of 10 children could justify (but not require) sacrificing one's own life, but it would be irrational to sacrifice one's life in order to ensure that one's niece receives an iPod.

Before we move on, it's worth pointing out that there is something counter-intuitive about the idea that the justifying and requiring force of reasons could be so mismatched. After all, it seems that if bringing about a certain state of affairs can justify so much, it must be because *that* state-of-affairs is very valuable. But if it is such a valuable state-of-affairs, why wouldn't it require some rather minimal sacrifices from us? And if it justifies more than bringing about another state-of-affairs it seems that this must be because it is a more valuable state-of-affairs. But wouldn't this consideration then require more? Of course, the more we are wedded to a teleological conception of practical rationality, the more counterintuitive this will sound. But this concern can be rephrased in a somewhat more inclusive manner. We can say that if a consideration justifies so much, this must be because it is awfully important. And if it justifies more than another consideration, this must be because it is more important, etc. Gert would reject this line of reasoning since it presupposes that reasons can be important only along one dimension. But it is important to note that, other than arguing that morality seems to give us an example of a similar structure (the fact that a certain course of action is very harmful to me can often make my not engaging in this course of

4 See *Brute Rationality*, especially 37-8.

action morally justified, but rarely, if ever, does it make it morally required that I not engage in this course of action),[5] Gert does not provide a 'direct' rationale for why reasons vary in strength along these two dimensions. He only shows that it is logically coherent to claim that there are these two dimensions, and argues that making this distinction allows us to accept a number of intuitive judgments about rational status. The word 'only' should not mislead us here; this is no doubt quite a lot. However, it is important to note that as long as other theories can do the same they are in an equal position. And if they propose a more intuitive kind of normative structure, they'll have an advantage over Gert's theory.

II

According to Gert, there are some judgments we want to make about irrationality that cannot be explained if we do not recognize that reasons can play these two different roles, and, in particular, if we do not recognize that the same reason can have different strengths in each of these roles. Gert appeals to at least two types of cases in order to illustrate and defend this claim. The first type involves cases of apparent intransitivity. (I'll call these selfless/selfish cases for simplicity.)

To illustrate this type of case recall the iPod/starving children example described earlier. In this example we need all the following to come out true:

> (a_1) It is permissible to face the risk of (near) certain death in order to save 10 children from death by starvation.
>
> (b_1) It is permissible to let 10 children die from starvation so as to save $200.

---

5 I find this comparison with morality unpersuasive. Although I can't go into this issue here, it might be worth quickly mentioning at least one problematic feature of this analogy. Even if one were to agree that ordinary morality has the requisite structure, moral reasons will only have this particular structure if it is true, as Gert supposes, that morality is concerned (almost) exclusively with our duties to others. But on this interpretation, morality by definition ignores any reasons that require the pursuit of our self-interest. So it is no surprise that our self-interest does not *morally* require that we act in any way. There is no similar feature of reasons in general that Gert could exploit in this context.

(c_1) It is not permissible to face the risk of (near) certain death in order to save $200.[6]

Taken together (a_1) — (c_1) generate an intransitive ordering of weights if we assume that these cases must be fully accounted for by precise weights reasons have along a single dimension. By (a_1) our reason to save 10 children is at least as strong as the reason to avoid the risk of death, by (b_1) our reason to save $200 is at least as strong as the reason save 10 children, but by (c_1) our reason to save $200 must be weaker than our reason to avoid the same risk of death. Of course, not all theories of rational choice accept all these claims. According to some subjectivist theories, for instance, there is nothing irrational to prefer (near) certain death over the loss of $200. However, as Gert correctly points out, these theories do not preserve intuitive judgments about irrationality. Even if we were not to use the fanciful word 'irrational' all that often, actions that express this kind of preference would be generally regarded as 'stupid,' 'dumb,' or 'insane.' One can also try to account for the intransitivity by suggesting that there is a measure of vagueness in these weights, so that each of the options in the above examples has a roughly equal, but not fully determinate weight. Although vagueness can often account for failures of transitivity, Gert has a persuasive argument to show that this case cannot be dealt with by treating the various options as roughly equal or by introducing some vagueness in the strength of reasons; no matter how vague one makes one's comparative judgments in each of these cases, vagueness alone cannot deliver us from an intransitive ordering. I'll not rehearse the argument here; I just want to say that I find it entirely persuasive if we assume that the options in (a_1)-(c_1) that appear to be tokens of the same relevant type are in fact tokens of the same relevant type. I'll challenge this assumption later, but for now I want to leave it in place. I'll grant for the time being that Gert has shown that one cannot account for this kind of case while accepting that there is only one normative dimension here. Thus, we'll accept for now that Gert has shown that these judgments of rational permissibility cannot be accommodated by a theory of rationality that assigns rational status to actions based solely on weights given to reasons along this single dimension.

The second type of case I'll call 'tiebreaker failures.' These cases are usually described as cases of 'incommensurability,' but I am introducing this unusual tag so as not to prejudice one's interpretation of the case.

---

6 Cf. examples on 22-3.

(a_2) It is permissible to prefer four hundred dollars over a month of suffering annoying (but not debilitating) pain.

(b_2) It is permissible to prefer not to suffer annoying pain for a month over four hundred dollars.

(c_2) It is permissible to be indifferent between not suffering annoying pain for a month and four hundred dollars

(d_2) It is permissible to be indifferent between not suffering annoying pain for a month and five hundred dollars.

Raz has proposed that the attitudes reflected in (a_2) — (d_2) (especially (c_2) and (d_2)) are best accounted for by assuming that the values in question are incommensurable. However, as Gert points out, this approach cannot explain the following intuitive judgments:

(e_2) It is not permissible to prefer fifty cents over not suffering annoying pain for a month.

(f_2) It is not permissible to prefer avoiding very mild and quick pain (say, taking a flu shot) over $1,000,000 (assuming you're not someone for whom this sum would not be very significant).

After all, if the values in question are incommensurable, how does it happen that as you make one of them very small or the other one very large, you end up with a very clear cut comparison? Again, I will not try to defend the proponent of incommensurability against this argument. In fact I am ready to grant Gert's point not provisionally, but definitively. I'll just assume in what follows that Raz's approach cannot account for the cases in question.

But it is important to note that once Gert shows that one approach fails for one kind of case, he does not consider how this approach would fare in the other kind of case. Gert proposes a unified solution for cases of tiebreaker failures and selfish/selfless cases. On the face of it, this is a theoretical advantage. After all, it seems better to account for multiple phenomena by a single principle than by multiple ones. But however theoretically laudable this is, the situation here is not that simple. First, on the face of it, these two cases are quite different. Cases of tiebreaker failures are puzzling cases of intrapersonal consistency. The fact that two different people have the preferences allowed by (c_2) and (d_2) would hardly present serious difficulties; a number of views, objectivist and subjectivist alike, can account for this difference in terms of differences in the circumstances of the different agents. On the other hand, the

selfless/selfish case is a puzzle of consistency of judgments of rationality across different people. Although Gert holds that (a_1)-(c_1) can all be true of the same agent, it is not clear that the same agent can *at the same time* be rationally disposed to behave in the ways allowed by (a_1)-(c_1).[7] Moreover, as mentioned earlier, the claim that there are two dimensions of normative strength does not have independent intuitive appeal. If different explanations have more independent appeal, the theoretical advantage would not lie so clearly on Gert's side.

III

Once we give up the idea that there must be a unified account of selfish/selfless cases and tiebreaker failures, the most obvious way of providing a compelling alternative to Gert's account is to invert the function of each of the accounts that Gert criticizes. That is, we can propose that vagueness accounts for the cases of tiebreaker failures and incommensurability (or something else with a similar structure) accounts for selfless/selfish cases. Let us start with tiebreaker failures. It's plausible to think that comparison across some values (for instance, pain and financial loss) has vague intervals. The vagueness can be understood in many different ways; it could be semantic, epistemic, or metaphysical.[8] But the important point is that in these cases an appeal to vagueness would deliver exactly the result we wanted; vague intervals would not imply that one cannot make clear-cut comparison when the differences are large enough, so the fact that it is vague whether $400 is better than a month of annoying pain does not imply that it is vague whether 50 cents is better than a month of annoying pain, or whether a million dollars is better than very mild pain.

Similarly, one can try to account for selfless/selfish cases by claiming that they present us with alternatives that are in some sense incommensurable. Although it is not often described in this way, Sidgwick's view about the comparative rationality of promoting one's own good as opposed to the impersonal good can be thought of as falling along these

---

7 I'll argue in more detail for this point later.

8 For different suggestions along these lines, see Thomas Hurka, *Perfectionism* (New York: Oxford University Press 1993), 87; Donald Regan, 'Value, Comparability, and Choice' in Ruth Chang, ed., *Incommensurability, Incomparability, and Practical Reason* (Cambridge, MA: Harvard University Press 1997); and Ruth Chang's Introduction in Chang, *Incommensurability*. Note that Chang's 'on a par' relation can also account for (e_2) and (f_2).

lines. According to Sidgwick,[9] there are two plausible but incompatible conceptions of rationality: according to the one conception, one should always promote one's greatest good; according to the other conception, one should always promote the greatest general good. However, according to Sidgwick, there is no decisive reason to accept one conception of rationality over the other. In its most extreme form, this view says that one can be irrational for choosing a lesser good to oneself over a greater good to oneself, or for choosing a lesser (general) good over a greater (general) good, but one can never be rationally criticized by how one ranks one's good in comparison to the general good. This kind of view fully explains the compatibility of (a_1)-(c_1). After all each of these claims is correct according to at least one live conception of rationality. And since we cannot settle among these different conceptions of rationality, we cannot say that any of the actions they allow are irrational.

A second strategy is to claim that altruistic reasons and self-interested reasons give rise to reasons at a different level, for a certain range of actions. According to this view, while reasons of immediate self-interest are primarily reasons to promote a certain state-of-affairs, altruistic reasons of a certain kind generate reasons to adopt general policies or dispositions. This view comes close to a Kantian understanding of imperfect duties, and duties of beneficence in particular. According to this Kantian view, for a certain range of actions, none of them, considered in isolation, is required by the principle of beneficence. However, when these actions are considered together, if we fail to perform any of them we may exhibit a pattern that violates the principle of beneficence.

Before evaluating how these two options compare to Gert's approach, it is important to notice an important difference between the two. The Sidgwick approach can account for virtually the same intuitions that Gert can since it permits, on the one hand, almost any amount of self-sacrifice,[10] and, on the other hand, the most unrestrained pursuit of self-interest. At the same time, the Sidgwickean view does not permit that I choose my own lesser good over my own greater good. The Kantian conception on the other hand cannot deliver the same result. Although the Kantian conception will be equally good at forbidding choosing lesser goods for oneself, it must render some callous actions irrational or at least against the demands of reason. Moreover the Kantian conception could not explain why it would be *for the same agent* rational to sacrifice life and limb and rational to refuse to donate a

---

9 I am not claiming great historical accuracy here.

10 'Virtually' and 'almost' are important qualifiers here.

modicum of money in choice situations that are *otherwise exactly the same*. So the Kantian account will probably not look better than Gert's unless one can show that there is really no intuition worth preserving here. But let us look more generally at how the strategies compare.

IV

Consider first the vagueness strategy with regard to tiebreaker failures. Vagueness seems to explain what goes on in these cases quite well. Moreover, this kind of approach seems to have greater intuitive appeal. We should probably expect either metaphysical or epistemic vagueness in our evaluative judgments; either the evaluative world is not very clearly delineated, or if it is, it is probably beyond our epistemic capacities (or at least would require Herculean efforts and astounding ingenuity) to pinpoint exactly where its lines lie. But vagueness approaches also have a further important advantage in pairing up options. It is possible that it is indeterminate that *a is P* and indeterminate that *b is P*, while it is not indeterminate that *if a is P then b is P*, or that *either a is P or b is P*. In order to see this, consider Michael who must pay a visit to the dentist. Let us assume that \$300 to \$600 is within the range of what it is permissible, at least for Michael, to pay for receiving optional anesthesia during his visit to the dentist. Let us now say that the following is true of Michael:

(a_v) Michael pays \$500[11] to be treated with anesthesia.

On both Gert's account and the vagueness approach, this turns out to be permissible. But let us assume that the following is true of Michael:

(b_v) Michael would not have paid \$400 to be treated with anesthesia.

(b_v) on its own is permissible on both accounts, and this is, of course, as it should be. However, it seems that (a_v) and (b_v) together indicate a subjectively irrational pair of attitudes. Now if one accepts the vagueness

11 Since, arguably, money is just an instrumental good, such considerations strictly speaking cannot be basic reasons in Gert's view. However given that it is so easy to compare monetary amounts, I'll keep using these examples. But one could substitute gastronomic pleasures, for instance, as long as one made sure that one of the pleasures was clearly greater than the other, and each great enough to justify the actions in question.

approach one has no problem accounting for the fact that (a_v) and (b_v) could be an irrational pair of attitudes, even if each of them is rationally permissible taken separately. It will be easiest to make the point by assuming that the vagueness in question is epistemic, but the same result can be delivered under each interpretation. The epistemic interpretation would claim that (a_v) and (b_v) are each separately permissible in light of the agent's epistemic situation: the agent does not know, within the $300-$600 range, whether it is better to forego the anesthesia or not. But a rational agent should know that if it is better to pay $500 to get the anesthesia, other things being equal, it would also be better to pay $400. So if the agent's best estimate of the evaluative situation in question is that he should pay $500 to get the anesthesia, he should conclude that his best estimate of the evaluative situation also requires him to pay $400 to get the anesthesia.

It is not clear how one could get this result on Gert's account. Saving $400 and saving $500 are both reasons that have enough justifying strength to make it rational that Michael decline anesthesia; neither of them has enough requiring strength to make it irrational that Michael purchase anesthesia. Presumably $500 has greater justifying and requiring strength than $400, but given that these differences in strength make no difference to the particular verdicts of rationality in Michael's situations when considered on their own, it's not clear how they could make a difference when considered in combination. Of course, one can add to the initial machinery to try to deliver this result. However one does this, one will need to distinguish the case in which we are comparing Michael's attitudes toward tradeoffs between pain and different amounts of money, and the case in which we are comparing Michael's attitudes toward tradeoffs between pain and a certain amount of money, or, say, pain and securing a certain aesthetic experience, or pain and securing a benefit to a stranger. However, even if this is possible, the advantages of theoretical unity have been lost; one has essentially conceded that cases of tiebreaker failure should get a different treatment from selfish/selfless cases. Moreover, in the absence of a specific proposal, it is not clear that the added machinery will not also appeal to notions of vagueness that we have proposed could account for the phenomena *without* relying on Gert's distinction between requiring and justifying reasons.

Now Gert might protest that this discussion does not take into account his theory of subjective rationality. Gert defends a reliabilist theory of subjective irrationality; actions that result from mental states that tend to lead to objectively irrational actions are subjectively irrational. Or in Gert's words:

> An action is subjectively irrational iff it proceeds from a state of the agent that (a) normally puts an agent at increased risk of performing objectively irrational actions,

> and (b) has an adverse effect by influencing the formation of intentions in the light of sensory evidence and beliefs.[12] (160)

However, it is not clear how this kind of approach could help. It says nothing about how we should treat combinations of attitudes as the ones described in (a_v) and (b_v). Moreover it is unclear how this kind of combination of attitudes, each of which would lead only to objectively rational actions, would put an agent in increased risk of irrationality. Finally, I am not sure that this account of subjective irrationality is an improvement over more standard accounts that take subjective irrationality to be understood in terms of objective irrationality relativized to the beliefs of the agent or the evidence available to him. Although discussing this issue in detail would lead us astray, I just want to point out an important shortcoming in Gert's analysis. Let us take, for instance, the thrill seeker. Having a disposition to seek thrills certainly puts you in an increased risk of performing irrational actions. However, the very same disposition might lead you to go skiing. Despite being led to ski by this disposition, you might still choose the same hills as other skiers of similar abilities, take the same precautions, etc. It is rather counterintuitive to say that when you are skiing you are being irrational simply because your action proceeds from a state that puts you at an increased risk of performing irrational actions, even when your choices, means, ends, and motives are indistinguishable from those of the rational skier.

V

Let us now look at our alternative approaches to selfish/selfless cases. I do not want to dwell too much on Sidgwick's approach, except to make a couple of suggestions. First Sidgwick's approach seems to do much better at explaining cases involving choices in which one makes a large sacrifice for someone else's benefit or fails to procure a great good for other people for the sake of procuring some relatively small good for oneself. At least some times these choices involve a great deal of regret and/or guilt, and Sidgwick can explain this fact by pointing out that given that these two principles of rationality are compelling, it is no surprise that one would have these negative feelings no matter what one chose. After all, in this case, one has chosen something that is irrational

---

12 Clause (b) is necessary in order to rule out causes of objective irrationality that are unconnected to the agent's will, such as 'blindness and clumsiness.' (*Brute Rationality*, 161).

according to a compelling conception of rationality. As for Gert's account, given that in these cases, one acted in accordance with a reason that had sufficient justifying strength to make the choice rational, it's not clear how such feelings could be warranted. Moreover, at least in one of the relatively few cases in which Sidgwick's and Gert's approaches disagree, it seems to me that Sidgwick's account gives the right answer. According to Sidgwick's approach, but not according to Gert's, the following combination of attitudes is rationally impermissible:

> (a_s) Larry donates $400 to UNICEF to save 10 children from starvation.
>
> (b_s) Larry would not have donated $400 to UNICEF if the money were to save 12 children from starvation instead.

Similarly to the case of (a_v) and (b_v), since both reasons can justify donating the money, and neither can require them, it is not clear on Gert's view what is wrong with this combination of attitudes. Notice that approaches that simply add agent-centered permissions to an agent-neutral calculus will have similar structures.[13]

Before we move on, it is worth noting that one may object to Sidgwick's account by claiming that it can make rational only the action that pursues the greatest good for oneself or the greatest good overall; any other choice would on Sidgwick's view be irrational.[14] This is no doubt a counterintuitive result; it seems perfectly rational to give some money to charity but not all the money that one could spare. This is a seemingly rational choice that, on the face of it, neither produces the greatest good for oneself (keeping all one's money) or the greatest good overall (giving

---

13 In fairness to Gert, he does examine a somewhat similar approach, namely Raz's claim that there are (or might be) exclusionary permissions. However, I find Gert's criticism of Raz unpersuasive, particularly if applied to any of these views. Very briefly, here are my misgivings. Gert asks about the 'ontological status' of exclusionary permissions or their 'justification.' I don't see why anyone should be more worried about the ontological status of any of the materials in these proposals than about the ontological status of two different kinds of strength that reasons may have. Moreover in terms of justification, as we pointed out above, there's no independent justification for the claim that the same reason has widely divergent requiring and justifying strengths; on the other hand, it is relatively easier to see the rationale for accepting that an agent-relative conception of rationality is as compelling as an agent-neutral one, or that there should be ways in which it is rational to restrict the claims that others' projects make on our own.

14 Gert raised this objection in 'Author Meets Critics' session on his book in the 2006 APA Pacific Division Meeting.

away all of one's money). However, the counterintuitive result is not implied by the Sidgwick strategy, at least the way I presented it. The strategy puts *no* constraints on the relative rankings of one's own good and the general good, so that it would be perfectly rational within this view to make various sacrifices for the greater good, but then refuse to make larger sacrifices to even greater goods. So suppose one could save 10 people from starving by donating $100, and 12 people by donating $1000. One might think that saving 10 people is worth a personal sacrifice of $100, but be unwilling to sacrifice $1000 even if this would result in more lives saved. This is perfectly rational on the Sidgwickean view presented here, but in this case one ends up choosing neither the option that maximizes one's own good (making no donations) nor the option that maximizes the general good (saving 12 lives).

VI

Let me now turn to the Kantian strategy. The following two claims make up, roughly, the Kantian view I have in mind:

> (I) We have non-derivative (requiring) reasons to promote the well-being (or interest) of others.
>
> (II) At least for a certain range of choice situations, we are not directly required to perform particular actions, but to adopt certain general policies of behaviour that apply to these particular actions.

To give a bit more substance to this sketch, let us consider the range of choice situations in which we have the opportunity to help others.[15] According to at least one version of this view, we are rationally required to help others, but not to the point that we are required to give their interests the same weight we give to ours. However, given the cumulative effects of helping others in situations that seem to demand very little from us, what we are required to do is not to attach a certain weight to their interests in any particular situation, but to act in accordance with general policies that are consistent with taking the well-being of others into account as reason requires. Violations of this rational requirement cannot be picked out by looking at particular occasions in which one fails

---

15 Although I am connecting this rationale to a Kantian view, it is not an exclusively Kantian one. With the exception of the emphasis on policies, it is quite similar to the rationale given by Scheffler to agent-centered permissions in *The Rejection of Consequentialism* (New York: Oxford University Press 1994).

to act in accordance with the requirement, but only by looking at a long stretch of the agent's life from which one can conclude that her actions are not in accordance with a permissible policy.

Note that it is compatible with this view to allow, for instance, that different agents (and even different agents at different times) are under different rational requirements depending on their desires, their subjective capacities, or dispositions. Moreover, with respect to the correct policy for an agent at a certain time, one can allow a range of permissible options, but one can also adopt a similar strategy to our earlier one; one can simply claim that within a certain range there is some kind of vagueness with respect to each agent. This would allow one to hold even the extreme view that for every agent there is an *exact* correct policy, and yet think that it is simply epistemically very difficult or impossible to do better than specify a range of policies that one has good reason to think approximate the correct one. This approach would allow that I keep $200 now, but give it on a different occasion, or even that I keep the money now but give larger sums on different occasions. Moreover, the theory can explain why the combination of attitudes in (a_s) and (b_s) strikes us as irrational. Whatever the details of the Kantian strategy, it certainly should not allow policies of helping others that are negatively sensitive to the number of people that can be helped.

But let us now look at a couple of things that the Kantian approach can't deliver. Notice that the Kantian approach seems unable to say that the *same* person could have all the attitudes in (a1) — (c1), or at least would not allow this to be the case if the person consistently fails to take up opportunities to use her money for charity. After all, how could the rationale given above justify a policy that would lead one to sacrifice one's life for the sake of others, but not to make a significant dent in one's comfort so as to save many lives? However, it is important to note that the view can deal with cases like those involving the (a1) — (c1) attitudes to a certain extent. It can, for instance, allow that our attitudes are different whether certain events are likely to happen or not, or whether we're in a particular relation to the person in question or not. So it might force me to make greater sacrifices in situations which are not likely to repeat themselves, or that were not likely to happen in the first place. What the Kantian strategy cannot allow, but Gert's approach can, is the following combination:

> (a_k) Larry faces the risk of (near) certain death to save 10 children from starvation.
>
> (b_k) Larry would not have paid $200 to save the same 10 children from starvation in the exact same situation.

That is, suppose Larry is ready to face this great sacrifice to save the children, but if it were to turn out at the last minute that he couldn't save the children this way, but that he could save them by paying $200 (or he could choose to pay $200 *instead of* risking life or limb), he would let the children die. In this case, according to the Kantian strategy Larry is irrational. But the Kantian strategy seems to be right about this, and, again, it seems to speak against Gert's approach that it has no clear way of showing that this combination is irrational. Gert might reply that Larry is effectively refusing to pay $200 to save life and limb and *that* is irrational on his account. However, Gert can't make that move and preserve the claim that (a1) — (c1) could be all rationally permissible, at least in the objective sense of rationality. For I take it that it makes no difference to our example if Larry was given an opportunity to save the children by donating the money at a different date. Suppose Larry has the following offer: he can pay $200 now to save the children or risk his life two days from now to save these same children (or, for that matter, other children who are just as unknown to him as these children). It seems that it would be just as irrational for him to say 'I want to keep $200 today; I'll save the children by risking my life tomorrow.' But this revised version of Larry's choices is identical to the ones allowed from (a1) — (c1) for a particular case in which the agent is the same in all cases, and has knowledge of the fact that she will face all these situations. In this case it seems hard to see how one can avoid moving from the subjective irrationality of the combination (a1) — (c1) for a fully informed agent to the objective irrationality of the combination.

Of course, if there is some significant difference between the children in (a1) or (b1) or if the agents in each situation are different, or even if the same agent just undergoes a change of heart from one choice to the other, (a1) — (c1) could come out to be true.[16] This is certainly what is behind our intuition that (a1) — (c1) are all true. However, accepting this point amounts to accepting that, whether or not we subscribe to the Kantian strategy, (a1) — (c1) are all true at the same time only if we hold either the relevant options or the agents as different in (a1) and (b1) respectively. Once we accept that, we have just supplied the resources that many of the traditional views of rationality need to explain how (a1) — (c1) can come out true without appealing to the distinction between justifying and requiring reasons.

---

16 Although in the last case the combination of attitudes would turn out to be rational only if either (a) the change of heart is a relevant change in the circumstances of the agent or (b) the sense of 'rationality' in question is the subjective one.

According to the Kantian strategy, behaviour that is callous enough in its general pattern so that one cannot see it as conforming to the principle of beneficence turns out to be irrational. The same goes, of course, for actions that violate the moral law independently of any general pattern of behaviour. Someone who simply shoves someone else out of the way to get to his car faster, or even perhaps, someone who foregoes the opportunity to save a child from drowning right in front of him in order to avoid dry-cleaning expenses, is, on this view, failing to act on a requiring reason. Gert finds this consequence of the Kantian view implausible. According to Gert, the Kantian must show that moral behaviour fails requirements of rationality that are compelling independently of one's acceptance of the norms of morality and whose inescapability does not presuppose any prior commitment to morality.[17] I must confess that I myself find it extremely implausible that the Kantian, or anyone else, can show any such thing. Although Gert is right that many Kantians seem to have this kind of ambition, at least one reasonably famous Kantian explicitly rejects it; namely, Kant.[18] The more plausible version of the Kantian strategy is that moral reasons are rational requirements that are not reducible to any other requirements or epistemically grounded on anything else; the immoral person is failing to act on reasons that cannot be derived either from more basic reasons or from thinner conceptions of rationality. One might complain that in this version of the Kantian strategy, the accusation of irrationality loses its force. But it's unclear what kind of 'force' one expects this

---

17 See *Brute Rationality*, 14. Gert only says that one has to show that requirements of rationality to which the Kantian appeals are inescapable, not that they are inescapable independently of prior commitment to morality. But I take it he would consider it trivial that they are inescapable in light of one's commitment to morality. Gert also says simply that Kantian is committed to show that the requirements can be described in nonmoral terms. But given that Gert thinks that requirements of rationality are such that it does not make sense to ask about them why we should comply with them, I take it that he would think that this implies the requirement I describe here.

18 Or at least the later Kant. A particularly striking passage is the following footnote in the *Religion*: 'The most rational being of the world might still need some incentives, coming to him from the objects of the inclinations, in order to determine his power or choice. He might apply the most rational reflection to these objects ... without thereby even suspecting the possibility of such a thing as the absolutely imperative moral law.... Were this law not given to us from within, no amount of subtle reasoning on our part would produce it or win our power of choice over to it.' (Immanuel Kant, *Religion Within the Boundaries of Mere Reason*, Allen Wood and George di Giovanni, eds. (New York: Cambridge University Press 1998), 27n.). It is more contentious whether the Groundwork is committed to such a project.

accusation to have; in this case it should not imply anything other than that the agent is failing to act in accordance with a valid normative requirement. It is true that we don't generally think about the crafty immoral agent as irrational, but, as Gert is well aware, there are various ways we can try to carve a subset of those actions that are appropriately called 'irrational' out of the set of actions that go against rational requirements in order to explain why it sounds odd to accuse Enron CEOs of irrationality. Gert argues that none of these ways delivers the right notion (or at least a very important notion) of rationality. Here I can only briefly raise one consideration that favours the Kantian 'wider' understanding of what can count of a rational requirement. Gert's position seems to have difficulty in accounting for certain 'deliberative phenomena' exactly because it proposes such a narrow conception of rational requirements. Suppose someone is trying to decide what she should do with some extra money she earned. From the perspective of the person deliberating, it seems that she is trying to weigh the various relevant considerations. She could invest the money in a retirement savings account, or send it to help her daughter who could do with a bit more money, or donate it to an emergency relief fund. In Gert's view, each of these reasons has enough justifying strength to make any of the possible courses of actions rational, since the accusation of irrationality would not stick no matter what the person chose. However, at least sometimes when we're trying to make such decisions, we think that we are trying to follow the weight of reasons (one could naturally say: 'I am trying to figure out what is best to do (or what I should do) with the money'). Moreover, when we deliberate, we think at least sometimes that making a mistake is a live option. So it seems to be part of our ordinary views that there are actions such that (a) they fail to follow the weight of reasons, and (b) it would not be appropriate, in ordinary discourse, to label them 'irrational.'[19] And I think it is the failure of Gert's view to make room for such cases that accounts for his failure to explain the irrationality of the various combinations of attitudes that we talked about earlier. After all, the problem in these cases is that the agent could not consistently think that she was following what reason requires in all these options.[20]

---

19 Notice that this cannot be explained in Gert's view by saying that some of these actions are subjectively rational but objectively irrational. After all, each of these actions should count for Gert as objectively rational.

20 Gert also appeals to the connection between irrationality and moral responsibility to bolster his case. But I must confess that I am not clear on the nature of the case here. (In fairness to Gert, he grants that the case is inconclusive. See *Brute Rationality*,

It is worth reiterating that I do not claim that any of these considerations are conclusive. I have just tried to show various roads available to the various extant theories of rationality that I think Gert has not adequately blocked. Without much further work, one cannot guarantee that they will not all turn out to be dead ends, and it is certainly one of the many virtues of Gert's outstanding book that it makes us see so clearly that further work is needed here.

Received September 2006

SERGIO TENENBAUM
University of Toronto
Toronto, ON M5S 1A2
Canada

82-3.) As Gert himself notes, not all cases of irrationality excuse or even attenuate one's responsibility. Many cases of irrationality are also cases in which the agent is not responsible, but there might be 'common causes' for the irrationality and the lack of responsibility (for instance, the agent is incapable of acting from the right reasons). Moreover actions done out hatred or revenge (and not for any kind of pleasure), or actions in which one puts others at risk by reckless endangering one's own life (reckless driving, or smoking heavily around one's children) could all (plausibly) come out irrational in Gert's account. And yet the agent bears the full responsibility of the harm that she may cause the others. I don't see how Gert can explain the attribution of responsibility in these cases if the connection between irrationality and moral responsibility is as tight as he thinks it is. Again, it is important to note that none of what I said in the defense of a wider conception of rational requirements depends on the details of the Kantian strategy; in fact, it is available to anyone who holds that moral reasons are (requiring) reasons *sans phrase*.