

*Reply to Tenenbaum*¹

JOSHUA GERT
Florida State University
Tallahassee, FL 32306
USA

In his clear, gracious, and helpful critical notice of *Brute Rationality*, Sergio Tenenbaum raises a number of worries about the way in which the central distinction in that book — the distinction between the requiring and justifying roles of practical reasons — is brought in to explain some commonplace normative data.² He also suggests two alternate ways of accounting for the same data, one of which derives from Sidgwick and has recently received its most sophisticated presentation in Derek Parfit's forthcoming *Climbing the Mountain*, and the other of which is inspired by Kant's distinction between perfect and imperfect duties. In this response to Tenenbaum I would like to do three things. First, of course, I would like to defend my own account. The cases Tenenbaum uses to argue against my use of the justifying/requiring distinction all have a common characteristic; roughly put, they are all ones in which it seems that some sort of maximizing principle applies, but is beyond the resources of my account. The same criticism has been raised in a number of other reviews.³ I would like to take this opportunity to make clear a general strategy that is available for dealing with all such cases. Next, I would like to show that both the Sidgwickian and Kantian al-

-
- 1 Thanks to Sergio Tenenbaum for comments on an earlier version of this reply.
 - 2 Sergio Tenenbaum, 'Brute Requirements: A Critical Notice of Joshua Gert's *Brute Rationality*,' *Canadian Journal of Philosophy* 37 (2007).
 - 3 Fred Schueler, *Mind* 115 (2006) 412-15; Terence Cuneo, *Ethics* 116 (2006) 785-7. Christian Miller and Paul Hurley also produced similar arguments in an author-meets-critics session at the Pacific APA in April of 2006.

ternatives either yield wrong conclusions or change the subject. And finally, I would like to show that both the Sidgwickian and Kantian alternatives in fact incorporate as robust a justifying/requiring distinction as I would wish to argue for.

I The Justifying/Requiring Distinction

Anyone reading this response to Tenenbaum's 'Brute Requirements' will already be familiar with the justifying/requiring distinction. But to avoid some possible confusions, and to facilitate later points, it may be worth making some quick remarks about that distinction. First, the distinction is not between two kinds of reasons, but between two roles that the very same reason might play. The distinctive thesis of *Brute Rationality* is best put as a thesis about the strengths of reasons in these roles: the thesis that many reasons have a great deal more justifying strength than requiring strength. Of course, since most views implicitly identify justifying and requiring strength, they have no need to distinguish the two roles. Consequently, it was important to distinguish and clarify them. But, despite a general lack of attention to the distinction, it is almost completely uncontroversial that reasons play the two roles, as I define them.⁴ For to play the requiring role is simply to do the following: to make it rationally required to perform some action that otherwise would not be rationally required. And to play the justifying role is simply to do the following: to make it rationally permissible to perform some action that otherwise would not be rationally permissible. Consider the reason 'that the action is necessary if I am to avoid significant pain.' This reason clearly can play the requiring role, since I can be required to perform some actions to which this reason applies, when, had the reason not applied, I would not have been required. Now consider the reason 'that the action will save a stranger from significant pain.' This reason clearly can play the justifying role, since it can be rationally permissible for me to perform some actions to which this reason applies, when, had the reason not applied, it would not have been rationally permissible. In light of these examples, it ought not be controversial that reasons can play both a justifying and a requiring role. Controversy enters at two subsequent points. First, I claim that (up to a certain degree of vagueness) we can assign any given basic reason degrees of strength in each of these roles: a justifying strength and a requiring strength. This should

4 'Normative Strength and the Balance of Reasons,' *Philosophical Review* 116 (2007) 533-62

not be overly controversial, since it amounts to little more than a denial of the controversial thesis of particularism about practical reasons. The real controversy enters at a second point: with the distinctive thesis of *Brute Rationality*, that justifying and requiring strength do not always covary. For example, an altruistic reason that involves the avoidance of significant harms for others can justifying a great deal more than a self-interested reason that involves the avoidance of less significant harms for the agent, even though the same altruistic reason might not be able to require as much as the same self-interested reason. Later in this response I will show that Tenenbaum's two proposals — the Sidgwickian and the Kantian — both entail this most controversial thesis. Our disagreement, therefore, is not about this, but about something else: about the explanation that lies behind the thesis.

II The Putative Counterexamples and a General Response

Tenenbaum's first putative counterexample involves Michael, who, quite rationally, decides to pay \$500 for optional anesthesia while having some dental work done. Tenenbaum stipulates unproblematically that Michael could just as rationally have decided to save the money, and suffer the pain and discomfort, and *also* that both choices would have been rational even had the anesthesia cost only \$400. My explanation of the rational permissibility of all of these choices makes use of one basic principle: whenever one of the options involved in an either/or choices is rationally permissible, this is because the justifying strength of the reason in favor of that choice is great enough to justify acting against a reason with the requiring strength of the opposing reason. Since reasons of roughly the same sort of normative significance, such as the reasons that appear in Michael's case, can have justifying strengths somewhat greater than their requiring strengths (that is, and to put it in a way that should make it seem very plausible: they can *justify* us in suffering more harms than they can *require* us to suffer), we get the pattern of judgments Tenenbaum stipulates in this case. So, for example, we can say that the reason 'this will help me avoid the pain of the dental work' can require someone like Michael to pay, say, \$50 (but no more), although it can justify him in paying up to, say \$700. Similarly, the reason 'this will save me \$400' can require someone like Michael to suffer the pain of, say, a pretty mean pinch (but no more) although it can justify him in suffering the pain of the dental procedure.

Tenenbaum's complaint about the general solution I offer for such cases is that it forces me to say that, supposing Michael had just chosen to pay \$500 for the anesthesia, he would be acting perfectly rationally if, immediately afterwards, upon being informed that there was an error,

and that it would cost only \$400, he refused it. That is, Tenenbaum's claim is that my solution entails that it is perfectly rational for Michael to refuse to spend \$400, but to be willing to spend \$500, for the very same thing. Tenenbaum also makes similar criticisms of my view using cases that involve both altruistic and self-interested reasons. In one such case, Larry donates \$400 to UNICEF to save 10 children, but would not have made the same donation, had he known it would save 12 children. And in another, more extreme case, Larry faces the risk of (near) certain death to save 10 children from starvation, but would not have paid \$200 to save the same 10 children from starvation in exactly the same situation. Tenenbaum holds that my view, as opposed to the views he suggests, requires us to say that these bizarre choices are all perfectly rational.

What Tenenbaum's criticisms omit, however, is adequate attention to the distinction between objective and subjective rationality. This distinction is as central to *Brute Rationality* as the justifying/requiring distinction. The objective rational status of an action is a matter entirely of the reasons that favor and oppose it, regardless of whether or not the agent is aware of those reasons, and regardless also of other options that the agent may have. In order to make this clearer, consider the following question, which I think would make sense, and have a clear answer, for most sane non-philosophers, whom I take to be the relevant authorities on first-order matters of rationality and value. The question is: does it make sense to risk breaking your leg in order to avoid burning to death in a house fire? Let us call this the 'break-your-leg' question. I think one answer that many sane non-philosophers would give to the break-your-leg question is a simple 'Yes.' To me this reveals something about the comparative normative significance of the risk of a broken leg and of burning to death. I would explain this fact about the comparative normative significance of these things in terms of the requiring and justifying strengths of the related reasons. Indeed, even someone who believes only in requiring strength might take essentially the same view. But once one starts thinking about Tenenbaum's purported counterexamples to my view, it can begin to seem that in some sense one cannot answer the break-your-leg question simply by reference to the strengths of the reasons that appear in the stipulated choice situation. This is because the existence of other options for saving one's life make a difference to the question of whether or not it makes sense to risk breaking one's leg to do so. Certainly it would be stupid to risk breaking one's leg to save one's life if one could just as well save it in a much less dangerous way. The existence of alternative options makes a difference, in an obvious way, to the question of whether or not an action of a certain description makes sense or not. So another answer to the break-your-leg question that seems pretty attractive at this point

isn't a simple 'Yes,' but a request for further clarification: 'What are my other options?'

Which is the correct response? My view is that there isn't a unique answer to this question. Rather, there are two ways of taking the break-your-leg question. One way of taking the question simply assumes that there are no other options. I think that this way of taking the question, and the fact that, taken in this way, almost everyone would give the same answer, reveals something about the comparative normative significance of the two harms involved in the question. Since I understand normative significance in terms of the justifying and requiring strength of the reasons such considerations provide, and also understand objective rational status in terms of justifying and requiring strength, I should be understood, when I ask about the objective rational status of an action, to be considering the action in a way that makes possible alternative actions irrelevant. My way of understanding objective rational status implies that to ask what the objective rational status of an action is, *given the existence of a certain alternative*, is to make a mistake.⁵ When I say that each of the options in Tenenbaum's putative counterexamples is rationally permissible, I should be understood to be speaking of the objective rational status. Objective rationality is a matter of the relative normative significance of the consequences of an action, and is not a matter of whether any given action that meets the description makes sense or not, regardless of context, etiology, or the existence of alternative options. By way of analogy, we can ask whether or not a certain piece of furniture is worth a certain amount of money or not. Even if the answer is 'yes,' it does not follow that if you want that piece of furniture enough to spend that amount, that it would always make sense to do so. For the same piece may be available for half the price across the street.

Understanding objective rationality in this way, I define subjective rationality in a reliabilist way: roughly, an action is subjectively irrational if it stems from a state of the agent that normally places that agent at increased risk of performing objectively irrational actions. This reliabilist account of subjective rationality makes it a very different beast than objective rationality. For one thing, subjective rationality will simply

5 My notion of subjective rationality takes account of an agent's subjective motivations and epistemic state, and for that reason is called 'subjective.' Nevertheless, there are some actions that it would be subjectively irrational for *any* agent to perform. There is a sense of 'objectively irrational' that could reasonably be applied to such actions, and that would be sensitive to the existence of alternatives. But that is simply not the sense at issue here, which has most directly to do with the relative values of items thought of as potential trade-offs.

not be a matter of the reasons favoring and disfavoring an action, even if we understand the relevant reasons in a subjective way, as, for example, beliefs that there are reasons. A good portion of *Brute Rationality* is devoted to explaining why this is so, and why almost all accounts of the relation between objective and subjective rationality go badly wrong by trying to make the two notions isomorphic in some way. One of the ways in which the two notions differ is that the subjective rationality of an action is easily affected by the existence of other possible choices.

Given these clarificatory remarks, I can explain why my view does not have any of the counterintuitive results that Tenenbaum, and many others, have supposed that it does. Consider the case of Michael and the optional anesthesia first. Let us first grant that the relative significance of the money and pain in the Michael example makes it true that it is *objectively* rational both to pay for, and to refuse to pay for, the anesthesia, whether it costs \$500 or \$400. This admission by no means commits me to the view that someone who has just agreed to pay \$500 could rationally then refuse the anesthesia when the doctor informs him that the price has been reduced by \$100. Subjective rationality is not simply objective rationality relative to the beliefs of the agent, as counterexamples of this general form seem to assume. Clearly a rational patient who is willing to pay \$500 for anesthesia cares a good deal about avoiding pain, and comparatively less about losing \$500. That is, this choice reflects the relative motivational strengths of desires to avoid pain, and desires to save money. Given such strengths, a refusal to pay \$400 must be explained. And the most likely explanations will involve the idea that something has gone wrong in the practical mental functioning of the agent, so that the refusal is almost certainly going to count as subjectively irrational on my view. For what sort of state could allow for this pattern of choice without placing the agent at increased risk of performing objectively irrational actions?⁶ And if the action proceeds from such a state, it counts as subjectively irrational on my account. To be clear, however, it does not seem to me *necessarily* irrational for someone who initially rationally agreed to pay \$500 to decide against the anesthesia moments later, even when the price is reduced. People do change their minds, and an account of rationality should allow for this. However, a patient cannot change his mind back and forth too many times without this very disposition to change one's mind about such matters

6 This question may prompt the response: 'a state that is stipulated to produce only actions that are objectively rational.' A full answer to this worry is beyond the scope of this short response, but depends on the fact that human beings are simply not subject to such states, and the idea that 'irrational' is best understood as a response-dependent notion.

becoming the very state in virtue of which his decisions will count, on my view, as subjectively irrational. So my account places limits on intra-personal inconsistency and fickleness, and not simply as stipulated additional conditions. I take this as an advantage. To put the point in general terms: virtually all of the formal constraints on preferences that decision-theorists require in order to prevent agents from becoming value-pumps — softened a to reasonable degree by the reference to the human norm implicit in my use of the phrase ‘increased risk’ — can be explained by appeal to the definition of subjective rationality. This shows that my view will also have no troubles capturing the irrationality of someone who donates \$400 to UNICEF to save 10 children from starvation, but would not have done so to save 12 children. And it also explains why, on my view, it would be irrational to refuse to pay \$200 to save ten children, *given* that one is willing to risk one’s life to do so. In all these cases we have to ask: what sort of state could explain this pattern of choice? If the answer is a state that places the agent at increased risk of performing objectively irrational actions, then my account will count the action as subjectively irrational.

III Two Other Solutions

1. *The Sidgwickian Proposal*

Now let us move on to Tenenbaum’s Sidgwickian explanation of our judgments in what he calls ‘selfish/selfless’ cases. Here is Tenenbaum’s example.

- a₁) It is permissible to face the risk of (near) certain death in order to save 10 children from death by starvation.
- b₁) It is permissible to allow 10 children die from starvation so as to save \$200.
- c₁) It is not permissible to face the risk of (near) certain death in order to save \$200.

All of these judgments seem correct both to me and to Tenenbaum, but if we try to account for them by assigning single strength values to the reasons provided by the various consequences, we get a striking failure of transitivity. (a₁) would entail that our reason to save 10 children is at least as strong as the reason to avoid the risk of (near) certain death. (b₁) would entail that our reason to save \$200 is at least as strong as the reason save 10 children. But (c₁) would require that our reason to save \$200 is weaker than our reason to avoid the risk of death. The strat-

egy for dealing with this presented in *Brute Rationality* should be clear: the relevant altruistic reason, like all altruistic reasons, has justifying strength far in excess of its requiring strength. That is, in this case, it can justify us in suffering the great harms in (a₁), but cannot require us even to suffer the minor harms in (b₁). Tenenbaum's alternative Sidgwickian account is the following. There are two competing conceptions of rationality: one makes the uniquely rational action that which maximizes personal good, while the other makes the uniquely rational action that which maximizes overall good. Sidgwick's own version of this view makes an action rationally permissible if and only if it is rational on one of these two conceptions. Thus Sidgwick's official account seems to give us only two rational options in any given choice situation: the one that maximizes impartial benefits, and the one that maximizes benefits for the agent. For example, given the following three choices

- 1) spend \$100 to save 10 people from starvation
- 2) spend \$1000 to save 12 people from starvation
- 3) spend nothing, and save no one

the official version seems to allow only (2) and (3), since — making some uncontroversial assumptions about the value of money and human lives from the impartial perspective — option (2) maximizes overall good, while (3) maximizes the agent's good. Tenenbaum and Parfit both modify Sidgwick in a way that opens up more rational options. In particular, Tenenbaum stipulates that there are no rational constraints on the relative importance one gives one's own good and the general good. That is, one can assign any ratio of relative importance to impartial reasons and to self-interested reasons, and if, on any such assignment, an action is favored by the balance of reasons, it is rationally permissible. This allows one to choose option (1). But importantly, it yields the conclusion that it would be irrational to

- 4) spend \$1000 to save 10 people from starvation

when one's other option was (1), or when one's other option was (2). Tenenbaum has claimed that in these latter respects, the Sidgwickian view is superior to the view put forward in *Brute Rationality*. Above I showed why this is not true. In general, the choice of (4) over (1), and the choice of (4) over (2), will come out as subjectively irrational on the reliabilist view put forward there, since, in general, the state that explains such a choice will be one that normally puts the agent at increased risk of performing objectively irrational actions.

Not only do the above cases not pose a problem for my account, but the way in which Tenenbaum's proposal explains their irrationality is too general, and ends up yielding verdicts of 'irrational' where we would want to withhold them. The way in which Tenenbaum explains the irrationality of choosing (4) over (1), for example, is the following. Suppose one values \$100 at h , and each person's life at v .⁷ If so, then while the value one gets from option (4) is $10v - 10h$, the value one gets from option (1) is $10v - h$. Option (4) is therefore never favored by the balance of reasons in the impartial perspective, no matter what the values of h and v are. For even more obvious reasons, (4) is not favored by the balance of reasons from the self-interested perspective, where v is no longer relevant. Thus, some elementary linear algebra yields that no matter what ratio of relative importance we assign to the impartial and self-interested perspective, (4) never is favored by the balance of reasons. Conclusion: it is not a rationally permissible choice. But now consider a choice situation in which one's options are the following:

- 5) give \$100 to an inferior charity because someone at work has asked one to do so, thereby saving one person from starvation
- 6) give the same \$100 to UNICEF, which one knows is much more efficient, and will save more lives with the money.

The reasoning Tenenbaum uses to give the correct result in the choice between (4) and (1) here operates to say that choosing (5) is irrational. True, there are some differences between (5) and (6) that go beyond the mere benefits to some unknown beneficiaries. For example, it will give some pleasure to one's co-worker to agree to donate, and it might also avoid some small social discomfort that would come from refusing. But it is a desperate expedient to claim that these small differences — the satisfaction of a co-worker, or one's discomfort at refusing the request — must be what make the difference here. Moreover, as long as we remain faithful to the Sidgwickian idea that significant normative indeterminacy is to be found only in the relations *between* the two perspectives, we can eliminate the possible relevance of such self-interested factors by assigning a subjectively acceptable monetary equivalent to the social discomfort, and adjusting the amount of money in (6) accordingly. Nevertheless, surely it is rationally permissible to choose (5). Surely it is rationally permissible to contribute to a charity, or to work for one,

7 Of course these simple assignments of value ignore such phenomena as the diminishing marginal value of money, and other factors. But none of these factors are important to the argument, so I am ignoring them.

even if one knows that there are other charities that are more efficient. Surely, for example, it is rationally permissible to volunteer at a public radio station, even if one knows that one's time could be spent in ways that benefit more people to a greater degree. Similarly, if a young child comes to your door peddling chocolates or girl-scout cookies, surely it is rationally permissible to buy them, even if one knows that the same money would be much more efficiently spent by Oxfam. On my view the restricted benefits such sub-optimal altruistic actions provide are enough to make them objectively rational. And since there is no reason to think that such sub-optimal choices must stem from a state of the agent that places the agent at increased risk of performing objectively irrational actions, they count as rationally permissible.

2. *The Kantian Proposal*

Tenebaum's second proposal for dealing with cases illustrated by judgments (a_1) through (c_1) is that '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.' The proposal is Kantian in spirit, in that it yields something like an imperfect rational requirement of beneficence. The idea is that one might ignore an altruistic reason in almost any particular case, but that altruistic reasons require one to form life-policies that give some real weight to the interests of others. Thus, in any given case, one is typically free to ignore the relevant altruistic reasons. For example one might spend \$200 on a new iPod for oneself rather than donating it to a charity that will use the money to save a number of children from starvation. However, someone who *always* made decisions of this sort, and *never* opted for the altruistic alternative, would reveal a life-policy that was — according to this proposal — irrational.

One problem I have with Tenenbaum's proposal is that it seems to me to give us the wrong answers. It seems to me that this happens in cases in which someone with sufficient power, and with the right sort of disposition, went through life quite happily ignoring the interests of virtually everyone except his own small circle of current friends. I think that there have been, and are, many people like this. Tenenbaum's view requires altruistic reasons to have a univocal (if somewhat vague) strength when we think of them as favoring and disfavoring various possible life-policies. And presumably he wants such policies to include, at least potentially, the risk of near-certain death for the sake of some unknown children. How it is that altruistic reasons with univocal strength values could render such policies rational at the same time they render it rational to behave, throughout one's life, in the selfish way described above, is very unclear. Of course one can try to avoid

this criticism by sticking to the claim that such a happy, selfish life really is irrational. But I cannot see what such a claim amounts to, or why we should be interested in an account of rationality that has this as a consequence. I am trying to capture a sense of irrationality that we are using when we call actions 'silly,' 'stupid,' 'crazy,' and so on, and that is closely related to mental illness, competence to give consent, and so on.⁸ The highly moralized notion of rationality that appears in Kantian accounts is ill-suited to this purpose, and does not even seem capable of drawing the essential distinction between irrational action and immoral action.

Of course, some Kantians will reject the distinction between irrationality and immorality. But there is a further and, I think, more destructive criticism of Tenenbaum's suggestion. When we think of altruistic reasons as justifying life-policies, we are no longer thinking of the sort of altruistic reasons that we take into account when deliberating, or when assessing the actions of others. When we are deliberating or assessing the actions of others the altruistic reasons of relevance have to do with the particular benefits and harms that other people will enjoy or suffer as the result of a particular choice: these are the kind of altruistic reasons with which *Brute Rationality* is concerned. But Tenenbaum's altruistic reasons are not of this sort. They justify life-policies. What, then, are these reasons? In (a₁) the relevant altruistic reason seems clearly to be the following: that ten children will be spared a painful death. But if one is justifying a life policy, these particular ten children are no longer of any direct relevance. Is the relevant altruistic reason the following: that there are an indefinite number of other people out there, whose welfare can be affected by one's attitudes and policies? Does it make sense to talk about the strength of *that* reason, in comparison with particular, context-specific self-interested reasons? When we talk about practical reasons, and debate the question of their strengths and whether or not they have two strength values or only one, the background picture is one according to which one can actually use the relevant reasons, and their strengths, to come to verdicts about particular actions, at least in some cases. Tenenbaum's Kantian suggestion is inconsistent with this picture.

8 Of course I do not mean that the aim is to capture a supposedly unique sense shared by precisely these English words. Certainly these are used in a variety of ways, some of which are quite distant from my notion of subjective rationality. Rather, the point is to capture the sense these words express when they express a certain important notion connected with competence, mental illness, freedom of the will, moral responsibility, and so on.

IV Are They Real Alternatives?

Recall the definition of the requiring role. A reason has strength in this role insofar as its presence can make the difference between an action's not being rationally required, and its being rationally required. Similarly, a reason has strength in the justifying role insofar as its presence can make the difference between an action's not being rationally permissible, and its being permissible. Consider, now, the Sidgwickian view, and the following fact about a certain action: the action will save a stranger from significant harm. This is the sort of fact that can be true of various different actions, and counts in favor of them: it is the sort of fact that we would intuitively wish to call 'a reason.' Can this reason play the requiring role? Well, given that there is no rationally required ratio according to which one must value selfish considerations in comparison to impartial considerations, this reason cannot require us to make even a trivial sacrifice. For there is an available ratio according to which the choice favored by the strongest reasons is the choice to not make that trivial sacrifice. On the other hand, there is also an available ratio according to which the action favored by the strongest reasons would be action on the altruistic reason, at least if the personal sacrifice was not, from the impartial perspective, even greater than the harm involved in the altruistic reason. This explains why it is rationally permissible to sacrifice one's life to save the life of a stranger, or to sacrifice one's leg to save the leg of a stranger, but not to sacrifice two of one's legs to save the leg of a stranger (at least, other things being equal). This shows that, on the Sidgwickian view as Tenenbaum and Parfit present it, altruistic reasons have considerable justifying strength, as I define that notion, but no requiring strength. The view therefore does not do away with the justifying/requiring distinction. Rather, it attempts to *explain* it. However, as I argued above, the explanation yields wrong conclusions: it does not allow the sort of suboptimal altruistic action that we see everywhere we look, and that we regard as perfectly rational.

Another way in which my view differs from the Sidgwickian view is that it allows for other sorts of reasons to have widely divergent justifying and requiring strength. For example, my view allows self-interested reasons that have to do with the gaining of benefits (as opposed to the avoidance of harms) to have considerable justifying strength, but little comparative requiring strength. Thus, one rational version of Robinson Crusoe might make great sacrifices for the sake of some theoretical knowledge, while another equally rational version of Crusoe might not be willing to make any significant sacrifice for the same knowledge. The potential of gaining theoretical knowledge provides a reason with very significant justifying strength, but little comparative requiring strength. Pleasure too provides reasons to Crusoe with far greater jus-

tifying strength than requiring strength. The Sidgwickian view has no way to account for this. For the way in which that view accounts for the justifying/requiring distinction depends crucially on a special feature of the impartial/self-interested distinction. That special feature is the following: any self-interested reason can also be viewed as an impartial reason. This is what prevents the view from yielding the conclusion that it would be rationally permissible to set oneself on fire in order to prevent some minor inconvenience to a stranger. The fiery sacrifice of one's life appears, in the impartial perspective, as the sacrifice of *someone's* life, and prevents it from being the case that there is some ratio of valuation of the two rational perspectives according to which the ridiculous altruistic act is favored by the strongest reasons. But the benefits — say, in terms of theoretical knowledge or pleasure — that can justify but not require certain personal sacrifices, may well be benefits for the agent. In such a case, the calculation one makes from the self-interested perspective will be the same as the calculation one makes from the impartial perspective. So the Sidgwickian view will have no way of accounting for the divergence of justifying and requiring strength for these reasons. Elsewhere I have also suggested that as self-interested reasons concern increasingly remote future times, their requiring strength diminishes, but their justifying strength does not. I cannot see how the Sidgwickian view can deal with this phenomenon either.⁹

Tenenbaum's Kantian proposal also does not actually deny the central thesis of *Brute Rationality* regarding the divergence of requiring and justifying strength for altruistic reasons. For to say that one is, in particular cases, free to give altruistic reasons almost any weight one wishes, is to say that, in particular cases, altruistic reasons have no significant requiring strength. And to say that it is rational to adopt a policy that allows one to act on altruistic reasons even in the face of very significant self-interested reasons is to say that, in particular cases, altruistic reasons have very significant justifying strength. Again, what Tenenbaum has done is offer an *explanation* for why this is so, in terms of a kind of meta-reason that is not the sort of reason we use to assess the rationality of particular actions, but that philosophers might use in order to assess the rationality of a life-plan. Tenenbaum's proposal does not explain *away* the justifying/requiring distinction. Rather, it simply tries to *explain* it. My main worries about Tenenbaum's proposal, therefore, are really worries about it as a potential explanation of the justifying/requiring distinction. One worry is that it is unnecessary: it posits, as

9 In fact, Sidgwick explicitly denies it. See *The Methods of Ethics* 7th ed. (London: Macmillan 1907; reissued 1962), 381.

an explanatory apparatus, a theory that itself stands in greater need of explanation than the phenomena it explains. Why, that is, do altruistic reasons have primary application only to life-policies, while self-interested ones apply primarily to particular actions? Tenenbaum's view, which makes altruistic and self-interested reasons differ in the *sorts of things* to which they primarily apply, seems to me to make them more wildly different than a view, such as mine, on which they differ merely in that altruistic reasons rationally require far less of us.

Of course, these two failed attempts to explain the disparity between the justifying and requiring strengths of altruistic reasons can lead one to wonder what really does explain the disparity. *Brute Rationality* does not offer any explanation of the sort Tenenbaum suggests. Rather, it takes the disparity as a brute fact about such reasons. In a similar way, the Sidgwickian view takes as brute facts the existence of two rationally distinct perspectives and the wide latitude of relative significance one can assign them. And the Kantian view takes as brute the primary applicability of altruistic reasons to life-policies, and the primary applicability of self-interested reasons to particular choices. From my perspective, the commitments of the Sidgwickian and Kantian views are more extravagant than my own, but it is certainly possible to disagree. Elsewhere I hope to go beyond this potential stalemate, and provide a rationale for the disparity, but that is a task for another book.

Received July 2007