

Imperfect Reasons and Rational Options

DOUGLAS W. PORTMORE

Arizona State University

Agents often face a choice of what to do. And it seems that, in most of these choice situations, the relevant reasons do not require performing some particular act, but instead permit performing any of numerous act alternatives. This is known as the *basic belief*. Below, I argue that the best explanation for the basic belief is not that the relevant reasons are incommensurable (Raz) or that their justifying strength exceeds the requiring strength of opposing reasons (Gert), but that they are *imperfect reasons*—reasons that do not support performing any particular act, but instead support choosing any of the numerous alternatives that would each achieve the same worthy end. In the process, I develop and defend a novel theory of objective rationality, arguing that it is superior to its two most notable rivals.

Agents often face a choice of what to do. And it seems that, in most of these choice situations, the relevant reasons do not require performing some particular act, but instead permit performing any of numerous act alternatives.¹ It seems, for instance, that I could now, in accordance with reason, do any of the following: watch TV, read a novel, practice the piano, volunteer for Oxfam, work on this paper, play with my daughter, or prepare for my next lecture. Following Joseph Raz, I will call this the *basic belief*. Raz calls it the *basic belief*, because it seems to be sufficiently well entrenched in our commonsense thinking that we should give it “credence unless it can be shown to be incoherent or inconsistent with some of our rightly entrenched views” (1999, 100). As he sees it, then, the basic belief is one of our starting points, and our task is to explain it rather than to defend it. However, even just explaining it is not an easy task, for it seems that there could be such rational options only if there were exactly equal reason to perform each of the optional act alternatives, and yet it is difficult to believe that such is the case.² For instance, it is difficult to believe that I have just as much reason to watch TV as to either volunteer for Oxfam or work on this paper. After all, volunteering for Oxfam seems vastly superior to watching TV in terms of the amount of impersonal good that it would do, and working on this paper seems vastly superior to watching TV in terms of the amount of good it would do me—assume that I’m not at this moment too tired to work productively on this paper.³ Moreover, the fact that, in many of these choice situations, the relevant act alternatives remain rationally optional even after there has been a slight increase

in the number and/or strength of the reasons that favor just one of the alternatives shows that their optional status could not have been due, in the first place, to a perfect balance of reasons (Gert 2008, 14). For instance, it seems that it would still be rationally permissible for me to continue to work on this paper even if there were a slight increase in the strength and/or number of the reasons that favor just one of the other options, as where, say, Oxfam institutes a new policy of giving each volunteer a delicious cookie. So the puzzle is to explain how, in most choice situations, there could be so many rationally optional act alternatives if, as seems to be the case, there is not exactly equal reason to perform each of them.

I will argue that this puzzle can be solved by taking a broader view of things and that, once we take this broader view, we find that there is in fact equal reason to do each of the more broadly conceived alternatives.⁴ As I will demonstrate, the relevant alternatives are not, as we've been supposing, the particular act alternatives available to the agent at a given moment, but rather the various courses of action that the agent could perform over the remainder of her life.

1. Imperfect Reasons and Rational Options

Often times, what an agent has most reason to do is determined by facts about what she has most reason to achieve. But such facts often fail to support any specific alternative, for there is often more than one way to achieve the same result. To illustrate, consider the following three examples. First, the fact that I need to get to the airport provides me with a reason to take any of the following equally attractive means to getting there: a bus, a taxi, or a train—assume that these are all equally attractive given their comparative cost, comfort, and convenience.⁵ Second, the fact both that I have a paper that needs to be finished in a month's time and that I will most likely finish it by then if and only if I implement a policy of spending two hours a day working on it provides me with a reason to spend two hours (any two hours) each day over the next month working on it. And, third, the fact that I need to spend one of the next two days grading exams and the other painting the backyard fence in order to meet certain important deadlines provides me with a reason either (1) to plan on spending tomorrow grading exams and the next day painting the backyard fence or (2) to plan on spending tomorrow painting the backyard fence and the next day grading exams.

These three sorts of reasons are what I call *imperfect reasons*, for they are analogous to imperfect duties in that they allow for significant leeway in how one chooses to comply with them.⁶ Just as the imperfect duty of beneficence

requires only that one be beneficent to a certain extent but leaves it up to one's discretion to whom and on which occasions to be beneficent, imperfect reasons speak in favor of achieving some worthy end but leave it up to one's discretion which of the equally attractive means to achieving that end to take.

Imperfect reasons arise in what Joshua Gert calls *multiple-option cases*, cases in which there is more than one equally attractive means to achieving the same worthy end (2003, 10). In multiple-option cases, it is rationally permissible to take any of the equally attractive means to achieving that end, for the relevant reasons are imperfect reasons, which do not support performing any particular act, but instead support performing any of the acts that would each achieve the same worthy end.

That there are multiple-option cases where the relevant reasons are imperfect reasons is, I take it, uncontroversial. But if the existence of multiple-option cases is as uncontroversial as I claim, we might wonder why philosophers such as Raz (1999) and Gert (2003) have held that we cannot account for the basic belief simply by appealing to them. The answer is that they each believe that multiple-option cases are insufficiently numerous to account for the basic belief.

To demonstrate this, Raz cites a case in which a woman named Mary has the opportunity to see a powerful performance of a good play at her local theater tonight but decides to stay home instead (1999, 99). Raz's view is that both alternatives—going to the theater and staying home—are rationally permissible. And although he admits that many reasons are imperfect reasons (or what he calls reasons that “are not time-specific”), he denies that such reasons allow us to account for the basic belief. He says:

In many cases the reasons for doing one thing or another are not time-specific: the same reasons and the same opportunity to conform to them will apply on a number, sometimes an indefinite number, of occasions. This does not, however, explain the basic belief. Quite apart from the fact that delay is not costless, the basic belief applies to time-specific reasons as well. Mary, in our example, does not have to go to the play even on the last evening of its run. She may still just not feel like it and do something else instead. (1999, 100)

Likewise, Gert believes that the set of cases in which there are rational options is significantly broader than the set of multiple-option cases. Gert believes, for instance, that it is rationally permissible both to sacrifice \$200 to prevent forty children from suffering from serious malnutrition for forty days and to refuse to do so, choosing instead to spend that money on oneself (2003, 8-9). As he sees it, this is a case where there is a rational option, but he denies that it is a multiple-option case. Gert says:

Multiple-option cases depend crucially on the fact that the class of rationally justified options in any given case can plausibly be seen as ways of doing the *same thing*. In the examples that were used to motivate the justifying/requiring distinction, on the other hand, the pairs of justified options were always of the following form.

- (a) Make a sacrifice for an altruistic reason.
- (b) Do not make the sacrifice.

Such pairs cannot plausibly be construed as alternate ways of pursuing the same end or acting on the same reasons [italicized letters are mine]. (2003, 13)

This, I will argue, is mistaken. I will show that options *a* and *b* can be construed as two alternative ways of pursuing the same end. As Gert admits, it is crucial to get the level of description right, for whether it's true that I am doing what I am rationally required to be doing depends on how we describe what it is that I am doing (2003, 13-14). To illustrate, consider again the case where I must get to the airport, and let us assume that I happen to be taking a taxi to get there. If we ask whether I'm rationally required to be doing *what I am doing*, the answer will depend on the level of description. If, on the one hand, we describe *what I am doing* as taking a taxi to the airport, the answer will be "No," because taking the bus is an equally attractive means of getting to the airport and I am rationally permitted to do that instead. If, on the other hand, we describe *what I am doing*, more generally, as going to the airport, the answer will be "Yes," because I am rationally required to go to the airport (or so we're supposing). So, at one level of description, I am rationally required to be doing *what I am doing*, and, at another, I'm not.

I think that the same can be said of Gert's example (and, as we'll see, of Raz's example as well). At a more general level of description, there are two courses of action that I might take over time, one that includes my acting self-interestedly now and altruistically later and another that includes my acting altruistically now and self-interestedly later, and it is plausible to construe these two courses of action as two ways of trying to achieve the same thing: a reasonably choice-worthy future.⁷ If we think that any reasonably choice-worthy future will contain both altruistic acts and self-interested acts, then it doesn't matter which of them I perform now and which of them I perform later so long as, in the end, I'll be performing a sufficient number of each.

In Raz's case, we can say something similar. If we acknowledge that Mary conceives of her agency as being extended over time, then we should think that, at the relevant level of description, the pertinent choice is not between seeing this play on its last night and staying home, but between two courses of action, containing both a certain amount of relaxation and a certain amount of cultural enrichment. So even though this is a good play and it's Mary's last opportunity

to see it, this will not be Mary's last opportunity to do something as entertaining and culturally enriching as seeing this play. And surely she isn't rationally required to take advantage of every such opportunity. On the broader view, then, what each of us, including Mary, should be trying to do is to complete a reasonably choice-worthy future, and any reasonably choice-worthy future will, for most of us, include both some relaxation and some cultural enrichment. It's important to note, then, that if Mary stays home tonight and relaxes, there will be other opportunities for her to similarly enrich herself, and that, if she goes to the play, there will be other opportunities for her to get the relaxation that she periodically needs. What's important, then, is not that she takes advantage of any specific opportunity for relaxation or cultural enrichment, but that she takes advantage of both sorts of opportunities sufficiently often over the course of her future.

So while admitting that there are multiple-option cases, Raz and Gert claim that such cases are insufficiently numerous to account for the basic belief, and so they resort to more contentious claims in trying to account for the basic belief: in Raz's case, to the claim that there are widespread incommensurabilities among reasons, and, in Gert's case, to the claim that reasons have two separable dimensions of strength (i.e., requiring strength and justifying strength). I, however, will argue that multiple-option cases are sufficiently numerous to account for the basic belief and, thus, that we needn't appeal to such contentious claims to account for the basic belief. My argument will be that Raz and Gert fail to appreciate the fact that most typical choice situations are multiple-option cases, because they fail to view our choices from the appropriate level of description: that of choosing between various future courses of action. To be clear, I am not going to argue that the contentious claims that Raz and Gert appeal to in accounting for the basic belief are false or even that they are not needed to account for something else. I will argue only that they are not needed to account for the basic belief and, thus, that we can account for the basic belief with only the theoretical apparatus that all sides agree on: imperfect reasons and multiple-option cases. Also, I should warn the reader that I will not be defending the basic belief. As Raz pointed out, our task is to explain it, and that's what I aim to do, and as simply as possible.

In the next section, I present my own theory of objective rationality, a theory according to which most choice situations are multiple-option cases. Then, in subsequent sections, I defend the theory.

2. The Future-Course-of-Action Theory of Objective Rationality

Before I state my theory, let me be clear on what it is a theory about. It is, as the name suggests, a theory about what it is objectively rational to do. Objective rationality is about what agents ought to do in the most fundamental and unqualified sense. So it's not about what they morally ought to do, nor is it about what they prudentially ought to do. Instead, it is about what they ought to do, all things considered. Thus, to say that an act is objectively *irrational* is to say that the "action absolutely should not be performed" (Gert 2004, 137). Some suggest that we should understand objective rationality in terms of advisability: objectively irrational actions are those that no fully-informed and well-meaning adviser would recommend performing on the assumption that the agent would take the advice.⁸ More precisely, though, we should understand the notion of objective rationality in terms of reasons: an act is objectively irrational if and only if there are decisive reasons for the agent to refrain from performing it.⁹ The objective rational status of an act is, then, purely a function of the reasons for and against performing it as opposed to its alternatives, regardless of whether or not the agent knows, or is even in a position to know, what they are.

By contrast, the subjective rational status of an act depends, not on what reasons there are, but on what reasons the agent takes there to be or, alternatively, on the practical mental functioning of the agent. On the former view, an act is subjectively irrational if and only if the agent has beliefs whose truth would give her decisive reasons to refrain from performing it. And when the agent has inconsistent beliefs, the act will be subjectively rational relative to some beliefs, but subjectively irrational relative to others (Parfit 2008). On the latter view, an act is subjectively irrational if and only if it indicates some failure in the practical mental functioning of the agent (Gert 2004, 160). I won't take a stand on which, if either, account of subjective rationality is correct. For my purposes, it is important only that we have a clear understanding of objective rationality, of which I now offer a theory.

At any given moment, t_i , prior to death, there are various ways a subject, S, might live out the rest of her life. For each possible way of completing S's life, there is a whole series of successive actions that S will perform if and only if S completes her life in exactly this way. Each of these possible ways of acting over the remainder of S's life is what I call a *future course of action*—an "FCA" for short—available to S at t_i .¹⁰ The theory of objective rationality that I endorse holds that whether or not it is rationally permissible (i.e., objectively rational) for an agent to perform a given act is a function of whether or not any of the available FCAs that include her performing this act are themselves rationally permissible.¹¹ Call this the *Future-Course-of-Action Theory* of objective rationality. More specifically, it says:

FCAT₁ S is, as of t_i , rationally permitted to perform a_i if and only if, and because, there is a rationally permissible FCA available to S at t_i in which S performs a_i .¹²

To better understand this theory, it will be helpful to consider a substantive version of it. The substantive version that I'm inclined to accept is *maximizing, teleological FCAT₁*:

MT-FCAT₁ (a) S is, as of t_i , rationally permitted to perform a_i if and only if, and because, there is a rationally permissible FCA available to S at t_i in which S performs a_i . (b) A specific future course of action, FCA_i , is rationally permissible if and only if, and because, it is, for S, maximally choice-worthy. (c) FCA_i is, for S, maximally choice-worthy if and only if, and because, there is no available alternative FCA that produces a life that S has more reason to desire than to desire the life that FCA_i would produce.

Clause *a* ensures that it's a version of FCAT₁. Clause *b* ensures that it's a maximizing version.¹³ And clause *c* ensures that it's a teleological version. Note that although MT-FCAT₁ is certainly a lot more substantive than FCAT₁, it is still quite ecumenical. There can be many different sorts of reasons for preferring one FCA to another, including both moral and prudential reasons. And although the theory is teleological, this doesn't mean that the grounds for preferring one FCA over another must always lie with the goodness of their respective outcomes; it could instead, or in addition, lie with the intrinsic features of the acts themselves, apart from whatever goodness these features possess.¹⁴ Thus, the proponent of MT-FCAT₁ could hold that, while pleasure is the only intrinsic good and pain the only intrinsic evil, an agent should prefer the life in which she refrains from murder to the life in which she commits one murder so as to prevent two others from committing murder even though the latter is the life in which she produces the most aggregate pleasure.¹⁵

3. MT-FCAT₁ and the Basic Belief

MT-FCAT₁ allows us to account for the fact that, in most choice situations, agents have numerous rational options. For instance, S will have the rational option to act either altruistically or self-interestedly at t_1 provided that there is both a maximally choice-worthy FCA available to S at t_0 in which S acts altruistically at t_1 as well as a maximally choice-worthy FCA available to S at t_0 in which S acts

self-interestedly at t_1 . Furthermore, S will have the rational option of performing any of numerous self-interested acts at t_1 (e.g., watching TV, working on a paper, going to the theater, etc.) provided that there is, in each case, a maximally choice-worthy FCA available to S at t_0 in which S performs that self-interested act at t_1 . And the same goes for S's performing any of numerous altruistic acts at t_1 . In this way, we can account for all the various sorts of acts that constitute the set of rationally optional alternatives. Since one FCA will often be just as choice-worthy as another if it contains the same proportion of various act-types, it simply does not matter which of these act-types I perform next so long as, in the end, the same balance among them is achieved. Thus, the basic belief only seemed puzzling, because we were taking too narrow a view of our choices. If we consider my choice of what to do from a perspective that ignores the fact that my agency is extended over time, it does seem difficult to believe that I have just as much reason to watch TV as to volunteer for Oxfam or to work on this paper. But when we take the broader view and compare FCAs (that is, whole series of actions that extend from now until my death) that all involve the same proportion of relaxing acts, altruistic acts, and career-furthering acts, just in different temporal sequences, we see that they are all equally well supported by reason, for they each constitute an equally effectual means to my completing a maximally choice-worthy FCA.¹⁶ Thus, I have the rational option of performing any one of them now.

To take a more concrete example, suppose that I am rationally permitted to perform either a_1 or a_4 at t_1 but that I am rationally required to perform one of these two at t_1 . This means that if we were to rank the FCAs available to me at t_0 , we would find (1) that there is, at least, one FCA in which I perform a_1 at t_1 that is maximally choice-worthy, (2) that there is, at least, one FCA in which I perform a_4 at t_1 that is maximally choice-worthy, and (3) that there is no FCA in which I perform neither a_1 nor a_4 at t_1 that is maximally choice-worthy. So my options are:

- (A) Undertake a maximally choice-worthy FCA by performing a_1 at t_1 .
- (B) Undertake a maximally choice-worthy FCA by performing a_4 at t_1 .
- (C) Undertake an FCA that is not maximally choice-worthy by performing, at t_1 , some act other than either a_1 or a_4 .

Note, then, that this is a multiple-option case, for there is more than one way to achieve the same end: specifically, that of completing a maximally choice-worthy FCA. And the reason that I have to perform a_1 at t_1 (viz., that doing so is a means to my completing a maximally choice-worthy FCA) is an imperfect reason, for it is equally a reason to perform a_4 at t_1 .

Return now to Gert's case, where the permissible options are:

- (a) Make a sacrifice for an altruistic reason.
- (b) Do not make the sacrifice [italicized letters are mine]. (2003, 13)

Gert denies that this is a multiple-option case, for he claims that “[s]uch pairs cannot plausibly be construed as alternate ways of pursuing the same end or acting on the same reasons” (2003, 13). However, he gives no argument for this claim; he must just think that it's obvious. It's not at all obvious, though. Indeed, as I've shown, options *a* and *b* could be two alternative ways of pursuing the end of completing a maximally choice-worthy FCA. Furthermore, options *a* and *b* can be construed as two alternative ways of acting on the same reason: specifically, the imperfect reason that I have to perform any act that constitutes a means to my completing a maximally choice-worthy FCA. Indeed, Gert's options *a* and *b* could just be my options A and B above.

Of course, if MT-FCAT_i is to account for the basic belief, it must do more than just account for options *a* and *b* being rationally optional, it must also account for the fact that they remain rationally optional even after the strength and/or number of reasons in favor of just one of them has increased slightly. Suppose, for instance, that I'm choosing between (a) spending today volunteering for Oxfam and (b) spending today watching football on TV. Assume that both are rationally optional, for let's suppose that it doesn't matter which I do today so long as I spend a sufficient amount of time doing each sort of activity over time. Let's further suppose that although Oxfam always provides each volunteer with a delicious cookie, today and today only it's going to be a freshly baked one. Here, then, there's a special reason for me to spend today as opposed to some later day volunteering for Oxfam: namely, I'll receive a more delicious cookie if I do.¹⁷ So the question arises: if, previously, there was a rational option for me to spend today either volunteering for Oxfam or watching football on TV when Oxfam wasn't offering a freshly baked cookie, why isn't the former now rationally required given that there is this new and additional reason for me to do so? Can MT-FCAT_i account for these two alternatives being rationally optional both before and after the addition of this new reason for my spending today volunteering for Oxfam?

Interestingly, MT-FCAT_i can provided that we accept the following two plausible assumptions, which I label 'A1' and 'A2':

- A1** If it is rationally permissible for an agent with options *a_i* and *a_j* to perform whichever one she presently prefers to perform (that is, whichever one she feels like performing), and if it is rationally

permissible for her present preference to shift from one option to the other, then it is true to say that she has, in the ordinary sense, the rational option of performing either a_i or a_j even though she is, in some stricter sense, rationally required to perform the preferred option.

- A2** Some slight increase in the number and/or strength of the reasons that favor the non-preferred option is typically insufficient to outweigh the strength of the reasons that the agent has to perform the preferred option.¹⁸

To illustrate the plausibility of A1 and A2, consider the following example. Suppose that Franklin is deciding whether to order soup or salad with his entrée. Assume that the two are equal in terms of cost, taste, calories, and nutritional content, but that Franklin just happens to prefer soup to salad on this particular occasion—perhaps, he just feels like having something hot. Even if we think that, given his current preference for soup, it would be objectively irrational for him to choose the salad, we still think that he has the rational option of choosing the salad instead, for we think that he is rationally permitted to come to have the opposite preference and that, were he to come to have the opposite preference, it would be rationally permissible for him to choose the salad.¹⁹ Given this, we should accept A1. And we should also accept A2, for even if we were to change the example so that the salad is slightly more tasty or nutritional than the soup, we wouldn't think that this slight increase in the strength of the reasons that favor the non-preferred option would be sufficient to outweigh the strength of the reasons that Franklin has to go with his preferred option.

Given A1 and A2, MT-FCAT_I can account for the fact that two options can remain rationally optional even after there is some slight increase in the strength and/or number of the reasons that favor just one of them. Consider, again, the case where I'm choosing whether to spend today volunteering for Oxfam or watching football on TV, and assume that all maximally choice-worthy FCAs include my spending today doing whichever one I prefer to be doing. So if, on the one hand, I prefer to watch TV today and volunteer for Oxfam on some later day, then all maximally choice-worthy FCAs include my watching TV today. But if, on the other hand, I prefer to volunteer for Oxfam today and watch TV (or do something else relaxing) on some later day, then all maximally choice-worthy FCAs include my volunteering for Oxfam today. Now let's suppose that, in fact, I prefer to watch TV today. In that case, all maximally choice-worthy FCAs include my watching TV today. And even if Oxfam offers the new enticement of a freshly baked cookie, it will still be the case that all maximally choice-worthy

FCAs include my watching TV today, for, as A2 states, this slight increase in the strength of the reasons that favor my volunteering for Oxfam today is insufficient to outweigh the strength of the reasons that I have to spend today doing what I prefer to be doing—namely, watching football on TV.

Of course, the reader may wonder: “In what sense is there a rational option to volunteer for Oxfam today if all the maximally choice-worthy FCAs include my watching TV instead?” The answer is: “in the ordinary sense.” As A1 states, if I’m rationally permitted to do whichever one I presently prefer to be doing, and if I’m rationally permitted to prefer either one, then it’s true to say that I have, in the ordinary sense, the rational option to do either. And, as I’ve just argued, these antecedents are true even on MT-FCAT_i. So even on MT-FCAT_i watching TV and volunteering for Oxfam are rationally optional both before and after the addition of this extra reason to volunteer for Oxfam today. It seems, therefore, that even a maximizing view such as MT-FCAT_i can account for the fact that two options can remain rationally optional (in the ordinary sense given by A1) even after there is a slight increase in the number and/or strength of the reasons that favor just one of them.

Things get a bit more interesting, however, when we up the stakes. Suppose, for instance, that Oprah Winfrey is offering, today and today only, to give to each person who volunteers today a brand new car. If so, it would seem foolish for me to pass up this unique opportunity to receive a free car, even if I would, other things being equal, prefer to watch TV today and volunteer for Oxfam on some later day. Here, then, there seems to be more than just a slight increase in the strength and/or number of the reasons that favor my volunteering for Oxfam today. And, thus, A2 doesn't seem to apply, which means that whatever reason I have to spend today doing what I prefer to be doing isn't going to be strong enough to outweigh the reason that I now have to spend it volunteering for Oxfam. So MT-FCAT_i would seem to imply that I am now, given Winfrey's offer, rationally required to spend today volunteering for Oxfam. And this seems to be the intuitively correct result.

But, now, consider a variation on this case, one in which Bill Gates is offering, today and today only, to donate an extra \$200 to Oxfam for each person who volunteers today. And let's assume that this extra \$200 would be enough to save an additional child or two.²⁰ It would seem that, here too, the new reason for volunteering today constitutes more than just a slight increase in the strength and/or number of reasons for doing so. So, again, A2 doesn't seem to apply. And, thus, MT-FCAT_i would seem to imply that I am rationally required to volunteer for Oxfam today even if I would prefer to just stay home and watch football on TV. This, I suspect, will seem counterintuitive, but let me suggest that MT-FCAT_i needn't have this implication, as I'll now explain.

As I think of it, MT-FCAT₁ has a rather sophisticated way of ranking lives and the FCAs that produce them. First, a full and meaningful life is to be preferred to most any life that is not full and meaningful—a full and meaningful life being one in which the agent adopts, and has a fair amount of success in achieving, a fairly broad range of ultimate ends.²¹ Second, of those available lives that are full and meaningful, the most choice-worthy ones are typically those that are best in terms of the agent's self-interest, where this, in turn, is often determined, in part, by which are the ones that best satisfy her preferences, including her present preference for doing one sort of activity rather than another. Third, although I don't have anything close to a complete account of what makes a life full and meaningful, it seems clear to me that an agent's life is not full and meaningful unless it exemplifies a life in which the welfare of others counts as being one of her ultimate ends. Taken together, these three claims imply that the strength of an agent's reasons to take advantage of some particular opportunity to benefit others will vary greatly depending on whether or not her doing so is essential to her leading a life in which the welfare of others counts as being one of her ultimate ends. If, on the one hand, her taking advantage of this particular opportunity is essential to her leading such a life, then she will likely have most reason to complete an FCA in which takes advantage of this opportunity. But if, on the other hand, her taking advantage of this particular opportunity isn't essential to her leading such a life, then she will likely have most reason to complete an FCA in which her self-interests are better served. And, in that case, A2 *will* apply, accounting for the fact that taking advantage of this particular opportunity is rationally optional as opposed to rationally required.

Here, we come close to what James Griffin calls a discontinuity in value. Take two values *A* and *B*. These values are *discontinuous* if "so long as we have enough of *B* any amount of *A* outranks any further amount of *B*" (1986, 85). So long as we have enough beneficence in our lives to count as leading a life that's full and meaningful, almost any life that's better in terms of our self-interest outranks a life that's better in terms of beneficence.²² Of course, we don't quite have a discontinuity in value here, for I don't think that the priority is absolute. It seems to me that even if one has enough beneficence in one's life to count as leading a life that's full and meaningful, a life that's exceedingly better in terms of beneficence could outrank even a life that's somewhat better in terms of one's self-interest. Likewise, a life that's moderately better in terms of beneficence could outrank a life that's only the slightest bit better in terms of one's self-interest.

So far, this has all been quite abstract, so let me try now to explain how this is supposed to work in more concrete terms. It seems to me, as well as to others

who have been similarly inspired by Kant's remarks on imperfect duties, that we are both morally and rationally required to adopt the welfare (or happiness) of others as an ultimate end. That is, we are required "to make the happiness of others a serious, major, continually relevant, life-shaping end" (Hill 2002, 206).²³ And this requirement in turn entails a further one: to act throughout our lives so as to exhibit a sufficient propensity to promote the happiness of others. Having such a propensity is, after all, constitutive of having adopted the happiness of others as an ultimate end. If, for instance, someone with the usual abilities and opportunities never does anything to promote the happiness of others, this would show that she has not truly adopted the happiness of others as an ultimate end (Hill 2002, 204). It's important to note, though, that having a propensity to *A* does not imply that one takes advantage of every favorable opportunity to *A*, as Robert Noggle explains:

[T]o say that Fred has a propensity to drink alcohol does not mean that Fred drinks whenever possible or even at every favorable opportunity. Fred may retain a propensity to drink even if he sometimes turns down drinks at parties that are good opportunities for drinking. Thus, one or a few failures to *A* even at good opportunities for *A*ing need not raise any doubts about a person having a propensity to *A*. However a pattern of continual failures to *A* despite favorable opportunities *would* cast doubt on the claim that a person has a propensity to *A*. (2009, 7)

Likewise, it seems that unless I have exhibited a pattern of failing to take advantage of good opportunities to help others, I can turn down this particular opportunity to help others by volunteering for Oxfam today and still count as having adopted the welfare of others as an ultimate end, and this is so even if this represents a particularly good opportunity for me to benefit others in that Gates is going to donate an extra \$200 if I do. So, in this case, it seems that I could lead a life that's full and meaningful whether I spend today volunteering for Oxfam or watching football on TV. My volunteering for Oxfam today simply isn't essential to my leading a life in which I count as having adopted the happiness of others as an ultimate end. And this means that the primary determiner of whether I should complete an FCA that includes my spending today volunteering for Oxfam or one that includes my spending today watching football on TV is which activity would best promote my self-interest, which in turn is likely determined by which I happen to prefer to be doing today.²⁴ So, despite initial appearances, A2 does apply to my situation—at least, assuming that I have not regularly failed to take advantage of good opportunities to promote the welfare of others.²⁵ And the applicability of A2 is what then explains the fact that spending today watching TV remains rationally permissible even after Gates provides me with an additional reason to instead volunteer for

Oxfam. Of course, it may be that were Gates to offer to donate enough extra money, it would, then, cease to make sense for me to spend today watching TV, for it may be that certain opportunities are so exceptionally good that one could not count as having adopted the welfare of others as an ultimate end and yet pass them up.²⁶ But the chance to ensure that Oxfam receives an additional \$200 doesn't seem to be this sort of "golden opportunity," as there are typically plenty of other opportunities to ensure that Oxfam receives an additional \$200.²⁷

These are, at any rate, my views on the matter. If readers have other views about how to rank lives and the FCAs that produce them, then they should plug those other views into MT-FCAT_i. And so readers should think that MT-FCAT_i has counterintuitive implications only if they think either (1) that there are instances in which an agent is rationally permitted to perform a certain act even though she ought to prefer some life in which she doesn't perform this act to any life in which she does or (2) that there are instances in which an agent is rationally prohibited from performing a certain act even though she ought to prefer some life in which she does perform this act to any life in which she doesn't. Myself, I can't imagine that there are any such instances. But even if others disagree, they should conclude only that MT-FCAT_i is problematic, not that FCAT_i is problematic. MT-FCAT_i is but one version of FCAT_i. Readers should think that FCAT_i has counterintuitive implications only if they think either (1) that there are instances in which an agent is rationally permitted to perform a certain act even though she is rationally prohibited from completing any FCA in which she performs that act or (2) that there are instances in which an agent is rationally prohibited from performing a certain act even though she is rationally permitted to complete some FCA in which she performs that act. So even if some doubt, for instance, that MT-FCAT_i can account for the full range of rational options that we take there to be, this doubt shouldn't necessarily carry over to FCAT_i, for there are satisficing versions of FCAT_i that could account for even more rational options than MT-FCAT_i can. Nonetheless, I believe, and have argued, that MT-FCAT_i *can* account for all the rational options that, upon reflection, we take there to be.²⁸ The point is only that even if others remain unconvinced, they shouldn't necessarily reject FCAT_i.

4. Accounting for the Rational Status of a Series of Actions

Like its leading rivals, FCAT_i can account for the basic belief. Unlike its leading rivals, FCAT_i can be easily adapted so that it also accounts for the rational statuses of series of actions. FCAT_i's rivals, by contrast, need more than mere modification; they need supplementation, and, as we'll see, this comes at a price.

Suppose that Smith enjoys gardening but doesn't enjoy getting his gardening tools out beforehand or putting them back afterwards. And suppose we're wondering whether it would be objectively rational for Smith to get his gardening tools out at t_5 and then, straight away, put them back at t_6 , leaving no time in between to garden. Here, we're wondering about the rational status of a series of actions. To evaluate such a series of actions, we need to modify FCAT_I, but only slightly:

FCAT S is, as of t_i , rationally permitted to perform A_i (where ' A_i ' is a variable that ranges over sets of actions) if and only if, and because, there is a rationally permissible FCA available to S at t_i in which S performs A_i .²⁹

Since a set of actions, A_i , can consist not only in a single act-token, a_i , but also in a series of successive act-tokens, α_i (let ' α_i ' be a variable that ranges over series of successive act-tokens), FCAT can account for the fact that although it is, as of t_1 , both rationally permissible for Smith to get his gardening tools out at t_5 (as where he's going to garden from t_6 – t_9) and rationally permissible for Smith to put his gardening tools back at t_6 (as where he has been gardening from t_2 – t_5), it is not, as of t_1 , rationally permissible for Smith to get his gardening tools out at t_5 and then, straight away, put them back at t_6 , for there is no rationally permissible FCA available to him at t_1 in which he performs this series of actions. There's just no point to Smith's getting his gardening tools out and then, straight away, putting them back, especially when there are other more enjoyable things that he could be doing.

We see, then, that FCAT_I can be easily modified (the result being FCAT) so that it can account for the rational status of a series of successive act-tokens—an *act-series*, as I'll call it. It's unclear, though, whether the same holds for FCAT_I's main rivals. Consider Gert's view. Gert accounts for rational options by supposing that, in most typical choice situations, many of the available act alternatives are supported by reasons that have more justifying strength than the opposing reasons have requiring strength.³⁰ For instance, Gert claims that the reason one has to benefit oneself has no requiring strength, but sufficient justifying strength to justify acting contrary to opposing reasons—at least, in most typical choice situations.³¹ In this way, Gert is able to account for the fact that agents have, in most choice situations, the rational option of performing any of various self-interested acts, such as going to see a movie, working to advance one's career, and spending time on one's hobbies. Gert also claims that, like the reason one has to benefit oneself, altruistic reasons have no requiring strength, but only considerable justifying strength.³² And, in this way, Gert is able to

account for the fact that agents have, in most choice situations, the rational option of doing either something altruistic or something self-interested. More precisely, *Gert's Theory* of objective rationality is this:

GT S is, as of t_i , rationally permitted to perform a_i if and only if, and because, it is not the case that S's performing a_i will bring some harm to the agent without bringing any compensating benefit to anyone (including, but not limited to, the agent). (Gert 2007a, 544).³³

Unfortunately, GT, unlike FCAT, offers no account of how to determine the rational status of an act-series. But, perhaps, we could rectify this by supplementing it in some way. But before considering how we might supplement it, we should look at FCAT's other main rival (viz., Raz's view), for it too is in need of supplementation.

Raz claims that, in most choice situations, the relevant reasons are incommensurable, and that, when two competing reasons are incommensurable, they neither defeat one another nor exactly balance out. Since whenever the competing reasons for various act alternatives fail to defeat each other, it accords with reason to perform any one of them, Raz (1999, 102-104) accounts for the basic belief by claiming that, in most choice situations, the relevant reasons are incommensurable. Unlike Gert, though, Raz never offers a substantive theory of objective rationality.³⁴ Sidgwick did, however, offer such a theory, and it's in the same spirit as Raz's view. Like Raz, Sidgwick accounts for rational options by supposing that the relevant reasons are incommensurable.³⁵ The only difference is that Raz goes much further than Sidgwick does, claiming not only that impartial reasons are incommensurate with self-interested reasons, but also that many different types of impartial reasons are incommensurate with each other and that many different types of self-interested reasons are incommensurate with each other. Thus, every act that is permissible on Sidgwick's view is also permissible on Raz's view; it's just that Raz thinks that there are a lot more permissible options than Sidgwick does. Given this, Raz's view will inherit any problem that Sidgwick's view has with respect to being too permissive. Since these are precisely the sorts of problems that I'll be concerned with in this section, it will suffice for my purposes to compare FCAT with *Sidgwick's Theory* of objective rationality:

ST S is, as of t_i , rationally permitted to perform a_i if and only if, and because, S's performing a_i will bring about either what's impartially best or what's best for S.

Like GT and unlike FCAT, ST offers no account of how to determine the rational status of an act-series. But, of course, we can rectify this by supplementing it in some way. There appear to be only three possible ways of supplementing GT and ST, adding one of the following three supplemental claims to each:

- S1** S is, as of t_i , rationally permitted to perform an act-series, α_i , if and only if, and because, each of the act-tokens within that act-series is, as of t_i , rationally permissible on the given theory.
- S2** S is, as of t_i , rationally permitted to perform an act-series, α_i , if and only if, and because, that act-series meets the same set of necessary and sufficient conditions for rational permissibility that individual act-tokens must, on the given theory, meet.
- S3** S is, as of t_i , rationally permitted to perform an act-series, α_i , if and only if, and because, there is a rationally permissible FCA available to S at t_i in which S performs α_i .

When we supplement GT with S1, S2, and S3, we get GT_1 , GT_2 , and GT_3 , respectively. Likewise, when we supplement ST with S1, S2, and S3, we get ST_1 , ST_2 , and ST_3 , respectively. Below, I show that no matter which supplemental claim they opt for the resulting theories are inferior to FCAT.

4.1. Supplementing with S1

The problem with both GT_1 and ST_1 is that they are implausibly committed to the following agglomeration principle for permissibility: If it is, as of t_i , both permissible for S to perform a_i and permissible for S to perform a_j , then it is, as of t_i , permissible for S to perform both a_i and a_j . To see why it's implausible, consider the following case, which I'll call the *Two Medicines Case*: Zeke is suffering from a potentially fatal disease. Fortunately for Zeke, he can cure himself either by taking medicine M1 at t_3 (call this a_1) or by taking medicine M2 at t_4 (call this a_2). He must be careful, though, not to do both, for taking these two medicines in succession is lethal.³⁶ GT and ST both plausibly imply that Zeke is, as of t_i , rationally permitted to perform a_1 as well as rationally permitted to perform a_2 . But, given this, both GT_1 and ST_1 further imply that Zeke is, as of t_i , rationally permitted to perform both a_1 and a_2 , and this is quite implausible. We should, therefore, reject both GT_1 and ST_1 . We should instead prefer FCAT, which avoids this counterintuitive implication given that there is no rationally permissible FCA in which Zeke performs both a_1 and a_2 .

4.2. *Supplementing with S2*

Given that both GT₁ and ST₁ are implausible, we might hope that GT₂ and ST₂ fare better, but such is not the case. They too have counterintuitive implications, for these views are unable to account for the fact that there are certain types of acts that are such that, although it is rationally permissible to take any individual future opportunity to instantiate that act-type, it is rationally impermissible to take every future opportunity to do so. To illustrate, consider what I'll call the *Hourly Wage Case*. Suppose that the highest paying job that I can get pays by the hour and allows me to work as many or as few hours a week as I choose. Assume that this is a job that I do not enjoy. Let's call all of the income that I have over and above what's necessary to buy what I need to continue to subsist and work my *surplus income*, and let's call all of the time that I have over and above that which I must spend eating, working, sleeping, etc. so as to continue to subsist and work my *free time*.

In this case, although it seems objectively rational for me to spend any particular hour of my free time working so as to earn more surplus income to send to Oxfam, it seems objectively irrational for me to spend *all* of my free time this way. If I were to spend *all* of my free time this way, I would thereby ruin my life as well as the lives of those that I most love. My marriage would eventually fall apart, my child would suffer greatly for the lack of my attention, and I would become severely depressed. And even though the good that would come from my working so many extra hours so as to send so much more money to Oxfam would more than compensate (impersonally speaking) for the havoc this wreaks on our three lives, I can't imagine that any fully-informed and well-meaning adviser would recommend that I act this way. Acting this way *might* be morally permissible, but it is certainly not rationally permissible. And yet both GT₂ and ST₂ imply that it's rationally permissible for me to act this way, spending all of my free time working so as to earn more surplus income to send to Oxfam.³⁷ GT₂ is committed to this, for even though I do incur certain harms in performing this act-series, these harms are more than compensated for by the good that my additional donations do.³⁸ And ST₂ is committed to this, for this act-series is, as we'll suppose, impartially best. FCAT, by contrast, is not committed to this, for it is plausible to suppose that, given my investment in, and commitment to, certain personal projects and relationships, there is no rationally permissible FCA in which I abandon them so as to perform this act-series.

4.3. *Supplementing with S3*

At this point, we're left with only GT₃ and ST₃. Because these two theories assess act-series in exactly the same way that FCAT does (after all, FCAT is equivalent

to FCAT₁ supplemented with S3), they all have the exact same implications with regard to the rational statuses of various act-series. This doesn't mean that they are all equally plausible, though. Regarding GT₃ and ST₃, we should ask: "Why, if the rational status of an act-series is to be assessed in terms of whether or not it is contained within some rationally permissible FCA, isn't the rational status of an act-token to be assessed in the same way?" Barring some reason either for thinking that act-tokens and act-series are to be assessed in different ways or for thinking that GT₃ or ST₃ can account for some intuitions that FCAT can't account for, we should prefer the more unified and systematic FCAT to the more fragmentary GT₃ and ST₃. But I can't think of any reason why act-tokens and act-series should be assessed differently. And I can't think of any intuition that GT₃ or ST₃ can account for that FCAT can't account for. FCAT is quite ecumenical. It can, as I've shown, account for the basic belief. And, like GT, it accounts for our intuition that it is objectively irrational to cause gratuitous harm to oneself.³⁹ Furthermore, it can account for our intuitions about various act-series, such as those with respect to the *Two Medicines Case* and the *Hourly Wage Case*. Indeed, there seem to be no intuitions that FCAT cannot account for. And if this is right, then the principle of parsimony directs us to supplant both GT and ST with FCAT rather than supplement either of them with S3. For FCAT is able to account for all the intuitive judgments that GT₃ and ST₃ can account for, only without having to make any additional, contentious claims, such as the claim that reasons for actions have two separable dimensions of strength or the claim that there is widespread incommensurability among our reasons for action.

5. GT's and ST's Counterintuitive Implications⁴⁰

In the last section, we learned that GT and ST are both in need of supplementation and that no matter how they're supplemented the resulting theories are inferior to FCAT. But this is not the least of their troubles. They have some very serious counterintuitive implications of their own, and this true regardless of how, or even whether, they are supplemented. Let's consider GT first. On GT, the objective rational status of an act does not depend on what the available alternatives are. This means, for instance, that whether it is objectively rational for an agent to sacrifice her life so as to save forty unidentifiable children from death by starvation does not depend on whether she could achieve the same result by instead sacrificing a mere \$200.⁴¹ This is quite implausible. Surely, whether it is rational to sacrifice one's life for the sake of achieving some worthy end depends on whether that end could be equally well achieved by making some more modest sacrifice. To illustrate, suppose that someone's \$200 cash donation to Oxfam is about to get ripped to shreds in some mail sorting machine

but that a nearby postal worker named Jane could prevent this from happening (thereby ensuring that forty unidentifiable children get the life-saving food and medical supplies that they need) by using her own body to jam the machine, resulting in her certain demise. Further suppose that if Jane refrains from throwing her body into the machine, then although this \$200 cash donation will be forever lost, she can continue to work for the rest of the day, thereby earning an additional \$200, which she could then send to Oxfam, saving forty unidentifiable children.

According to GT, it is objectively rational for Jane to throw her body into the machine, thereby sacrificing her life so as to ensure that the envelope containing \$200 in cash makes its way to Oxfam.⁴² But this is clearly objectively *irrational* given that Jane could just as easily ensure that Oxfam gets an extra \$200 by instead continuing to live and sending Oxfam a check for \$200.⁴³ FCAT avoids such counterintuitive implications, because, on FCAT, the objective rational status of an act does depend on what the available alternatives are, and given that it's possible to save forty unidentifiable children by merely sacrificing \$200, there is no rationally permissible FCA available to Jane in which she instead sacrifices her life to achieve the same end.

In reply to such counterexamples, which I base off of Sergio Tenenbaum's counterexamples (2007), Gert says:

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. 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. (2007b, 467)

But Gert's way of understanding objective rationality is exactly the same as my way of understanding objective rationality, which I spelled out at the beginning of Section 2. An objectively irrational act is one that absolutely should not be performed (2003, 137); it is one that no fully-informed and well-meaning advisor would recommend (2003, 140)—at least, not on the assumption that complete advice will be given and all of it will be taken.⁴⁴ On this understanding, it seems perfectly clear that the objective rational status of an action does depend on the existence of certain relevant alternatives. For instance, whether it is objectively

rational to spend \$500 on a particular piece of furniture doesn't just depend on whether that piece of furniture is worth \$500, but also on whether that same piece of furniture is available for half that price across the street, for no fully-informed and well-meaning advisor would recommend spending \$500 dollars on a piece of furniture at one store when it could be got for half that at the store across the street. Likewise, whether it is objectively rational for Jane to sacrifice her life to save forty unidentifiable children depends on whether or not she could do so by instead sacrificing a mere \$200.

Gert's claim that the "[o]bjective rationality is a matter of the relative normative significance of the consequences of an action, and is not a matter of...the existence of alternative options" is simply false. The question of whether or not Jane's life is less normatively significant than the lives of forty children is not at all the same question as whether or not Jane should, in her given circumstances, sacrifice her life to save forty children. And objective rationality concerns the latter, not the former. Objective rationality is about what one should and should not do; it's not about the comparative normative significance of certain harms or benefits. So Gert's reply to these counterexamples is completely inadequate. These counterexamples give us reason to reject GT, and the fact that FCAT is not subject to these same counterexamples is a significant advantage for the theory.

GT, therefore, already has one huge strike against it, but it has other counterintuitive implications as well, some of which it shares with ST. For instance, both GT and ST imply that some morally obligatory acts are objectively irrational.⁴⁵ For instance, it is, on these two theories, objectively irrational to do what one must do to keep a promise if it involves some non-trivial risk of harm to oneself for which there is no compensating benefit to another. To illustrate, suppose that I promised my landlord that I would give her a \$100 cashier's check, and let's suppose that, given the associated bank fees, it's going to cost me \$105 to purchase one. Further suppose that we both need the money and will, therefore, each incur some non-trivial harm if we're out any amount of money. And let's assume that the harm that I'll incur if I spend \$105 to purchase the cashier's check is slightly greater than the harm my landlord will incur if she doesn't get the \$100 that I owe her. According to both GT and ST, then, it is objectively irrational for me to spend \$105 to keep my promise. In doing so, I would incur a non-trivial harm for which there is no compensating benefit to another, and my doing so would maximize neither my utility nor impartial utility. I suspect, though, that few would regard such dutiful conduct as objectively irrational.⁴⁶ I did, after all, promise my landlord the check, and it's wrong for me to break this promise even if my keeping it is slightly worse for me than my breaking it is for her.

Both theories also yield counterintuitive results in the following sort of case. Suppose that there is a burning building with twelve children trapped inside—call this the *Burning Building Case*.⁴⁷ Assume that I've called the fire department, but that it is obvious that they won't arrive soon enough, and so all the children will perish unless I act now and rescue them myself. Unfortunately, I lack the proper fire-fighting equipment, and so there is no way for me to rescue the children without suffering severe injuries in the process. Moreover, assume that the more trips that I make into the burning building the more severe my injuries will be. And assume that I can rescue, at most, two children per trip into the burning building, one under each arm. Given the great costs associated with each trip, let's assume that what would be best for me, self-interestedly speaking, is for me to refrain from making any trips, saving no children. But let's also assume that what would be best, impartially speaking, is for me to make six trips, saving all twelve children.

ST implies that it is objectively rational both to save all of the children and to save none of the children, but objectively irrational to save any number of children in between. GT implies that it is objectively rational to save any number of children from zero to twelve, including any odd number of children. Both implications are counterintuitive. Contrary to ST, it seems objectively rational to act so as to strike some sort of compromise, saving some of the children but not making so many trips as to incur the most severe injuries. And, contrary to GT, it seems objectively irrational to save any odd number of children given that, in each case, there is no additional cost to carrying out, on a given trip, two children as opposed to just one.⁴⁸ For instance, if I'm going to take a third trip into the burning building and suffer the associated injuries, I should carry out two children, not one, saving a total of six, not five. But, on GT, the reason that I have to carry out two children as opposed to just one has no requiring strength at all. Thus, on GT, it is objectively rational for me to carry out only one child even though I could have saved an additional child at no extra cost to myself or anyone else. I can't see how anyone could regard such an action as objectively rational, recommending it to someone.

FCAT, by contrast, avoids these counterintuitive implications—or, at least, it does given certain plausible assumptions about what does and doesn't count as a rationally permissible FCA. First, FCAT can account for the fact that it is objectively rational to perform morally obligatory acts even when such acts involve some harm to oneself for which there is no compensating benefit to another, for it is plausible to suppose that there are some rationally permissible FCAs that will include such acts. After all, the rational permissibility of an FCA is a function of moral reasons as well as non-moral reasons, and there is a strong moral reason to fulfill one's moral obligations. Second, FCAT can account for the

fact that, in the *Burning Building Case*, it may be rational to strike some sort of compromise and rescue some even number of children between zero and twelve.⁴⁹ There is, after all, no reason to assume that the only rationally permissible FCAs are those in which I rescue either no children or all twelve children. Third, FCAT can account for the fact that, in the *Burning Building Case*, it would objectively irrational to rescue any odd number of children, for it is implausible to suppose that there is any rationally permissible FCA in which I save any odd number of children given that there would be no reason for me to choose such an FCA over one that included my saving an additional child at no extra cost to me or anyone else.

6. Objections to MT-FCAT

The particular version of FCAT that I favor is *maximizing, teleological FCAT*:

MT-FCAT (a) S is, as of t_i , rationally permitted to perform some set of actions, A_i , if and only if, and because, there is a rationally permissible FCA available to S at t_i in which S performs A_i . (b) A specific future course of action, FCA_i , is rationally permissible if and only if, and because, it is, for S, maximally choice-worthy. (c) FCA_i is, for S, maximally choice-worthy if and only if, and because, there is no available alternative FCA that produces a life that S has more reason to desire than to desire the life that FCA_i would produce.

I've argued that MT-FCAT can account for the basic belief and that it avoids the sorts of problems that plague both GT and ST. Of course, some may think that MT-FCAT has problems of its own. So, in this section, I attempt to rebut various possible objections to MT-FCAT. Some of these objections are specific to MT-FCAT and others apply more generally to FCAT. I'll start with the latter.

6.1. Does FCAT implausibly presuppose that agents can form the intention to complete particular FCAs?

Proponents of FCAT can, and should, deny that agents can form the intention to complete particular FCAs—at least, they should if we're talking about human agents who have a significant amount of time left to live. Such agents have long and complex FCAs to choose from, and such FCAs are beyond a human's capacity for comprehension. Certainly, no human could ever form the intention to perform each and every act contained within such a long and complex FCA. A human could, however, form the intention to achieve an end that requires her to

perform a very long and complex series of actions. To illustrate, consider that, when I was younger, I formed the intention to earn a Ph.D., and this was an end that I could achieve only by performing a very long and complex series of actions. At the time, I had only the vaguest notion of all that I would need to do. So even though I did not form the intention to perform some highly-specified series of actions culminating in my receiving a Ph.D., I did form the vague intention to do pretty much whatever it takes (within reason) to earn a Ph.D. The idea that agents can form these sorts of intentions is not at all problematic, and this is all that the proponent of FCAT needs. An agent can form the intention to complete her life in reasonably choice-worthy way without knowing specifically what all that entails.

6.2. *Is it problematic that FCAT holds that whether it is objectively rational for an agent to perform a given act depends on what she could do were she to perform that act, not on what she would do were she to perform that act?*

Suppose that I promised my friend Dave that I would attend his wedding as his best man. As it turns out, if I show up to the wedding as promised, I'll end up drinking too much, become belligerent, and pick a fight, physically assaulting one of the other guests. Unbeknownst to me or anyone else, this is what I *would* do were I to show up. Of course, this is not to say that I couldn't show up, drink moderately, and behave myself. I *could*. Indeed, let's assume that I have the sort of control over my drinking that compatibilists believe is essential for moral responsibility. It's just that I don't always exercise this control. Sometimes, as in this particular instance, I choose to drink more even when I know that it's imprudent for me to do so. Now, my showing up, drinking moderately, and behaving myself is what's best for all concerned. Second best is my breaking my promise by not showing up. Worst of all is my showing up, drinking too much, and physically assaulting one of the other guests.

There are at least two schools of thought on whether I should show up as promised. On the one hand, there's *actualism*, the view that an agent ought to φ if and only if what *would* happen were she to φ is better than what *would* happen were she to refrain from φ -ing. And, on the other hand, there's *possibilism*, the view that an agent ought to φ if and only if what *could* happen were she to φ is better than what *could* happen were she to refrain from φ -ing. So, whereas the possibilist holds that I ought to show up, the actualist denies this.⁵⁰ As should be clear, FCAT is a version of possibilism, and, as such, it holds that I should show up as promised even though what will happen if I do is going to be worse than what would happen if I don't.⁵¹ To some, this seems problematic.⁵² To others, it seems more problematic to deny this, for how can we consistently claim, as the actualist does, that I ought not to show up as promised when, as we all admit, I

ought to show up, drink moderately, and behave myself.⁵³ After all, I can't show up, drink moderately, and behave myself if I don't show up. Of course, actualists will claim that possibilism has implications that are at least as counterintuitive as this one, and I, unfortunately, don't have space here to address these worries. Instead, I'll settle for making the following two points.

First, there's still a lot of new work that's being done on the relative merits of these two schools of thought and so it would be premature at this point to dismiss FCAT simply on the grounds that it's a form of possibilism. Second, even if we were convinced that possibilism was false, we should not dismiss the overall thrust of the paper, for it seems possible to develop a non-possibilist version of FCAT that would still be able to account for the basic belief. I don't have space here to show specifically how that might be done, but it is something that I've demonstrated elsewhere.⁵⁴

6.3. Does MT-FCAT implausibly suppose that the best way to balance the various competing factors in our lives is going to be the same for each of us?

We all have to balance various competing factors in our lives: e.g., career, family, entertainment, personal projects, and the welfare of others. But it would be a mistake to assume that the best way to balance these competing factors is going to be the same for each of us. It's not as if the best life for each of us is the one in which certain proportions are achieved, such as where, say, 40% of our resources goes to our careers, another 40% goes to our family and friends, and the remaining 20% goes to our helping others. As I see it, there are no ideal proportions. Instead, what's important is that each of us leads a life that's full and meaningful. And this means that we should each adopt certain ultimate ends—the welfare of others being one of them. But the fact that we should each have this end doesn't entail that we should each dedicate some preset percentage of our resources to helping others. Indeed, someone with abundant resources and limited familial obligations will need to dedicate a much higher percentage of her resources to helping others than, say, someone else with limited resources and a disabled child to care for. And not only are we not required to dedicate the same percentage of our resources to a given end, we're not even required to adopt the same set of ends. Even if there are certain broadly defined ends that we are each required to adopt, these ends will, in practice, yield different specific ends for different people. Take the end of developing one's talents. For one person, it will mean adopting the development of one's musical gifts as an end. For another, it will mean adopting the development of one's philosophical talents as an end. So, different people will need to take different paths to a full and meaningful life. Moreover, of those paths that a given person could take to a full and meaningful life, some will be more choice-worthy than others given

one's individual psychological traits. Consider, for instance, that although we all need enjoyment in our lives, we enjoy different things. Whereas some find skydiving exhilarating, others find it terrifying.

In sum, the best way to balance various competing factors in our lives is going to vary from person to person, for a lot is going to depend of each person's individual circumstances and psychological traits. This is true both when it comes to determining which sorts of FCAs result in a full and meaningful life and when it comes to determining which of these FCAs are the maximally choice-worthy ones.

6.4. Does MT-FCAT implausibly suppose that, for each of us, there is only one way to optimally balance the various competing factors in our lives?

As noted above, how any individual should optimally balance the various completing factors in her life is going to depend a great deal on her own individual circumstances and psychological traits. But even given a particular set of circumstances and psychological traits, there can be more than one way to optimally balance two or more competing factors. For instance, trading \$30,000 for a Honda Accord might, for a given individual, be just as good as trading \$70,000 for a Jaguar XJ. Each of these two trades may strike, what is for her, a maximally good trade off between savings and luxury. Likewise, there may be more than one way for a person to strike an optimal balance between, say, altruism and self-interest. Imagine, for instance, a young doctor who must choose to work either for The Mount Sinai Medical Center in New York or for Doctors without Borders in Africa. The former offers the chance to live in greater material comfort and to employ more sophisticated medical techniques, whereas the latter trades less of these goods for the greater opportunity to save lives. Both, then, involve a certain tradeoff between altruism and self-interest, but both tradeoffs may be just as good. So, on MT-FCAT, it may be that, for each person, there is a range of different tradeoffs among various competing factors that all count as optimal.

6.5. Is it problematic that MT-FCAT implies that it can be rationally permissible for an agent to perform each of the acts in an act-series that is itself rationally impermissible?

As Warren Quinn has shown, certain theories of rationality, such as MT-FCAT, imply that it can be rationally permissible for an agent to perform each of the acts in an act-series that is itself rationally impermissible. To illustrate, Quinn cites of the following sort of case:

Suppose there is a medical device that enables doctors to apply electric current to the body in increments so tiny that the patient cannot feel them. The device has 1001

settings: 0 (off) and 1 . . . 1000. Suppose someone (call him the self-torturer) agrees to have the device, in some conveniently portable form, attached to him in return for the following conditions: The device is initially set at 0. At the start of each week he is allowed a period of free experimentation in which he may try out and compare different settings, after which the dial is returned to its previous position. At any other time, he has only two options -- to stay put or to advance the dial one setting. But he may advance only one step each week, and he may *never* retreat. *At each advance he gets \$10,000.*

Since the self-torturer cannot feel any difference in comfort between adjacent settings, he appears to have a clear and repeatable reason to increase the voltage each week. The trouble is that there *are* noticeable differences in comfort between settings that are sufficiently far apart. Indeed, if he keeps advancing, he can see that he will eventually reach settings that will be so painful that he would then gladly relinquish his fortune and return to 0.

The self-torturer is not alone in his predicament. Most of us are like him in one way or another. We like to eat but also care about our appearance. Just one more bite will give us pleasure and won't make us look fatter; but very many bites will. And there may be similar connections between puffs of pleasant smoking and lung cancer, or between pleasurable moments of idleness and wasted lives. (1990, 79)

MT-FCAT implies that it is rationally permissible for the self-torturer to advance a step each week, for, each week, there is a maximally choice-worthy FCA available to him in which he advances that step. Yet the net result over time is that the self-torturer ends up completing a series of actions that is itself objectively irrational, for completing such an act-series results in his being worse off than he would have been had he never advanced the dial at all.

Is this problematic? I'm not so sure. Clearly, it's lamentable. The fact that one might find oneself in a situation where one could act rationally in every instance and yet end up leading a life that no rational person would choose to lead is indeed lamentable and in the same way that collective action problems are lamentable. With regard to collective action, it's lamentable that sometimes we find ourselves in situations where if each of us acts rationally in pursuit of our own self-interest, we will all end up worse off than we would have been had we all acted less rationally. And because such situations are lamentable, it is rational for us to do whatever we can to avoid ending up in such situations—e.g., by entering into Ulysses contracts that constrain our future actions and by allowing governments to establish incentives for our cooperating with one another. But to say that such situations are lamentable is not to say that those theories that imply that such situations can arise are problematic. And even if such theories were problematic, it's clear that this wouldn't be a problem for FCAT in particular. GT and ST also imply that it is rationally permissible for the self-torturer to advance a step each week. Indeed, as Warren Quinn has shown, any theory that holds

that agents ought, at every moment, to do the best they can will have this implication.

Myself, I find it hard to deny (at least, at times) that it is rational for the self-torturer to advance a step each week, and yet I acknowledge that this will, over time, result in the self-torturer being worse off than he would have been had he not advanced at all. Is this a lamentable situation to be in? Yes. Does it follow that any theory that holds that it's rational to advance a step each week is mistaken? I don't think so.⁵⁵ Of course, we should reject any theory that implies that it is rationally permissible for a person to perform the act-series that the self-torturer ends up performing, as theories such as GT₁ and ST₁ do. But, fortunately, MT-FCAT correctly implies that performing such an act-series is rationally impermissible.

6.6. Is it problematic that MT-FCAT implies that as we get close to the end of life our rational options narrow considerably?

Suppose that I'm deciding whether to spend this Sunday at the beach with my family or at home alone watching football on TV. On MT-FCAT, both would typically be rationally optional, for there would typically be both some maximally choice-worthy FCAs available to me in which I spend today watching TV and some later day doing something with my family as well as some maximally choice-worthy FCAs available to me in which I spend today doing something with my family and some later day watching TV (or doing something else similarly relaxing). What's important is that I take advantage of both opportunities to relax and opportunities to spend time with my family and that I do each sufficiently often, but it doesn't generally matter which I do on any particular occasion. But now suppose this is my last day of life. In that case, I have no more future opportunities. And so the choice is no longer between doing one of these activities today and the other on some later day, but is instead between doing one of them today and the other never again. Given this, I tend to think that my current preferences are not going to be decisive in determining which option is maximally choice-worthy, even though they might have been decisive had the choice been between which to do today and which to do on some later day. I'm inclined to think that if my choice is between spending my last day with my family and spending it alone watching TV, I ought spend it with my family even if I happen to prefer to spend the day alone watching TV, for I'm inclined to think that I *ought* to prefer the life in which I spend my last day with my family to the one in which I spend my last day alone watching TV. And, if this is right, MT-FCAT implies that I am rationally required to spend the day with my family. But far from being objectionable, this implication seems to be a merit of the view. After all, I can't see how any fully informed and well-

meaning advisor would recommend that I spend my last day alone watching TV. Surely, there is something better I could be doing on my last day! So I think that MT-FCAT is right to imply that as we approach the end of our lives our rational options narrow considerably.

But now consider a different sort of case. Suppose that I'm on my death bed dictating my last will and testament. And let's assume that I've already provided for my family's continual financial security, as I have taken out a rather substantial life-insurance policy with them as the beneficiaries. Some might claim that, having provided for my family's financial security, I am rationally permitted to do almost anything with my estate that would not be harmful, leaving it to any individual or organization or even using it to purchase some extravagant funeral for myself—one in which, say, my body is blasted into outer space (see Cochrane 2009). I'm not so sure. First, let's take the case where I have so far exhibited throughout my life a sufficient concern for others, having clearly adopted the welfare of others as one of my ultimate ends. Since having such a concern for others doesn't necessitate taking advantage of every good opportunity to help others, I could pass up this opportunity to help others and still count as having led a full and meaningful life. Given this, it seems that it might be rationally permissible for me to use the money for my own benefit by, say, creating an open-access, digital archive of my work. Of course, it might also be permissible for me to donate my entire estate to Oxfam. Which of these two options would in the end be most choice-worthy would depend, I think, on which would best serve my self-interests, which in turn would depend on whether I happen to care more about helping others to an even greater extent than I already have or more about ensuring that my work is made widely available. I don't, however, think that it would be rationally permissible for me to leave my estate to some random stranger or to spend it on some extravagant funeral in which my body is blasted into outer space. Given my particular cares and concerns, either action would be positively foolish.

Now, consider the case where I haven't lived a life in which I've exhibited having adopted the welfare of others as one of my ultimate ends. Unfortunately, nothing I can do now will change this. I was required to have adopted the happiness of others as "a serious, major, continually relevant, life-shaping end" (Hill 2002, 206), and I didn't. The best I can do now is to atone for this moral failure. So the question becomes, "Which life is more choice-worthy: the one in which I atone for my previous selfishness by leaving most, if not all, of my estate to an organization such as Oxfam or the one in which I continue to act for my own benefit, using the money to set up a digital archive of my work?" I think that the former is more choice-worthy. And if this is right, then MT-FCAT implies that it would be rationally impermissible for me to use the money for my

own benefit. Of course, all this is based on the assumption that a life in which I use the money for my own benefit is not as choice-worthy as some of other life that's still available to me. But if the reader thinks otherwise, then she should think that MT-FCAT has different implications. For instance, if it turns out that I'm wrong and that the life in which I use the money to set up the digital archive of my work is maximally choice-worthy, then MT-FCAT will imply that using my estate to do this is rationally permissible. Again, the reader should think that MT-FCAT has counterintuitive implications only if she thinks that it is rationally permissible (or impermissible) for an agent to do what is incompatible (or compatible) with her leading a life that is maximally choice-worthy. Furthermore, the reader should think that FCAT has counterintuitive implications only if she thinks that it is rationally permissible (or impermissible) for an agent to do what is incompatible (or compatible) with her completing a rationally permissible FCA.

7. Conclusion

I've argued that the best explanation for the basic belief is not that the relevant reasons are incommensurable or that their justifying strength exceeds the requiring strength of opposing reasons, but that they are imperfect reasons. I've argued that, ultimately, what agents have most reason to do is to complete an FCA that is maximally choice-worthy. The reason an agent has to perform a particular act that constitutes a means to completing a maximally choice-worthy FCA is an imperfect reason, for it equally supports all of the other particular acts that also constitute a means to completing such an FCA. Since, in most choice situations, many of the available act alternatives constitute a means to completing a maximally choice-worthy FCA and, thus, a way of acting on this imperfect reason, we find that the relevant reasons permit performing any of these various act alternatives. And this, I've argued, is how we best account for the basic belief. In the process of making these arguments, I have developed and defended a new theory of objective rationality, which I call FCAT. I've argued that this theory is promising not only because it enables us to account for the basic belief, but also because it has a number of distinct advantages over its two main rivals, such as being able to avoid certain counterintuitive implications and to offer a more unified and systematic account of the rational statuses of both act-tokens and act-series.⁵⁶

Notes

¹ An act alternative is any of the mutually exclusive and jointly exhaustive act-tokens that are available to a given agent at a given time.

² In this paper, I will be concerned with objective rationality, not subjective rationality. An act is objectively irrational if and only if the agent has decisive reason to refrain from performing it, and objectively rational otherwise. The objective rational status of an act is purely a function of the reasons for performing it as opposed to its alternatives, regardless of whether or not the agent is aware of these reasons. By contrast, the subjective rational status of an act depends, not on what reasons there are, but on what reasons the agent takes there to be, or, alternatively, on the practical mental functioning of the agent—see Parfit (2008) and Gert (2004), respectively. I will have more to say about this distinction in Section 2 below. The reader should assume that, unless otherwise qualified, I am talking about objective rationality whenever I use the word ‘rational’ or any of its variants.

³ Indeed, it seems that we *know* that watching TV isn’t the alternative that I have most reason to perform. Thus, our belief that the relevant reasons permit my watching TV can’t be explained in terms of any kind of epistemic uncertainty as to whether or not it’s the alternative that I have most reason to perform.

⁴ There are other potential solutions to this puzzle, including: (1) adopt a satisficing conception of rationality, wherein it is rationally permissible to perform any act that is supported by sufficiently weighty reasons, and then claim that, in most choice situations, the reasons that support the various act alternatives are all sufficiently weighty (Slote 1989); (2) claim both that reasons for action have two separable dimensions of strength and that, in most choice situations, the reasons that support many of the act alternatives all have sufficient justifying strength to justify acting against the opposing reasons with whatever requiring strength they have (Gert 2004); and (3) claim that, in most choice situations, the relevant reasons are incommensurable such that it is rationally permissible to perform each of the incommensurable alternatives (Raz 1999, 102-104). In Sections 4 and 5, I provide some reasons for preferring my proposed solution to such rivals.

⁵ I borrow this example from GERT 2003, p. 10.

⁶ Jonathan Dancy calls such reasons “unfocused reasons,” and Raz calls a subclass of such reasons “reasons that are not time-specific”—see Dancy (2004, 100) and Raz (1999, 100). I first introduced the notion of an imperfect reason as well as the idea that it could be employed to account for rational options in Portmore (2000)—see pp. 199-201.

⁷ I talk about a reasonably choice-worthy *future*, because we can make choices only with regard to how the future goes and not with regard to how the past went. That said, one should not be misled into thinking that I’m committed to the view that the past is irrelevant in determining which possible futures are reasonably choice-worthy. For instance, it may be, as some have argued, that the value of a life composed of two or more temporal parts standing in certain relations to one another need not equal the sum of the values that those temporal parts would have if they existed alone, apart from those relations. Some have argued, for instance, that how much some misfortune suffered at t_1 detracts from the value of one’s life as a whole depends on whether or not it will be redeemed by events that occur later in life—see, for instance, Portmore (2007) and Velleman (1993). So suppose that it is now t_9 and I have the choice of performing either a_1 or a_2 , where a_1 will, and a_2 will not, redeem the misfortune that I suffered at t_1 . We can certainly hold that the future in which I perform a_1 is more choice-worthy than the one in which I perform a_2 and on the grounds that, if I undertake the former as opposed to the latter, my life will be better *on the whole*.

‘Reasonably choice-worthy’ can be given either a maximizing or a satisficing interpretation. On the maximizing interpretation, a future is reasonably choice-worthy if and only if it is maximally

choice-worthy. On the satisficing interpretation, a future is reasonably choice-worthy if and only if it is sufficiently choice-worthy.

⁸ See Gert (2004, 140) for more on this idea. I add the qualification “on the assumption that the agent will take the advice” to eliminate the worry that in some cases the agent might be predisposed to rebel against such advice. I should also add that the advice needs to be complete. That is, it won’t do to advise an agent to do x without also advising her to do y when doing x without doing y would be disastrous. What’s more, we must realize that an agent’s reasons for doing x could potentially diverge from an advisor’s reasons for advising her to do x , and the objective rationality of an agent’s doing x is a function of her reasons for doing x , not the advisor’s reasons for advising her to do x . At best, then, the advice model is just a sometimes useful (but also sometimes misleading) heuristic for getting a grip on our intuitions about what is objectively rational. It must, however, be used cautiously. For more on this, see Portmore (2009a).

⁹ Another helpful way of understanding this notion, which was suggested to me by Peter de Marneffe, is as follows: an act is objectively irrational if and only if the set of true propositions provide decisive reasons for the agent to refrain from performing it. I’m assuming, here, that a reason for S to perform x is just some fact that counts in favor of S ’s performing x . Thus, there can be reasons for S to perform x of which S is completely unaware, for S can be completely unaware that the relevant reason-providing facts obtain.

¹⁰ What I call a “future course of action” (an FCA) is sometimes called a “maximal sequence of acts”—see, for instance, Carlson (1995, 121-123) and Goldman (1978). Accordingly, we can understand an FCA as a maximal sequence (or series) of actions, where an act-series, α_i , is maximal if and only if there is no act-series available to the agent of which α_i is a proper part. At the very end of life, the available FCAs will each contain only one last act.

¹¹ I acknowledge that FCAs are typically far too complex to be the objects of our conscious choice. But this doesn’t mean that there aren’t reasons to choose to complete one FCA as opposed to another. The fact that Smith would, other things being equal, lead a better life if he were to become a professional philosopher is a reason for him to choose to complete an FCA in which he becomes a professional philosopher. And this is precisely the sort of reason that might enter into his deliberations when deciding whether or not to enroll in a Ph.D. program in philosophy.

¹² This theory is adapted from the sort of theory of unconditional moral obligation that Feldman (1986) and Zimmerman (1996; 2006a) defend. I call it “FCAT_i” (the ‘ T ’ standing for ‘initial’), because it is only my initial formulation. Although, in the abstract above, I refer to it as ‘a novel theory of objective rationality’, this *kind* of theory is not at all new. What’s new is only that I’m applying Feldman’s idea to give an account of objective rationality.

A few points of clarification: First, I use the word ‘because’ to mean ‘in virtue of the fact that’. Second, I use the word ‘available’ to mean the same thing it does for those who define act-utilitarianism as the view that an act is morally permissible if and only if there is no *available* act alternative that would produce more utility than it would—see Carlson (1995, chap. 5) and Zimmerman (1996, p. 46) for more precise accounts of this notion, which is also sometimes called *personal possibility*. Third, given that ‘ought’ implies ‘can’ and that what one is able to do can change over time, it is important to state the theory in terms of what is, as of a certain time, permissible—see Feldman (1986, esp. 11-12), Portmore (2009a, chap. 7), and especially Goldman (1976) for further explanation. Fourth, the variable, a_i , ranges over intentional actions. Since forming an intention to perform a given act is not itself an intentional action, FCAT is silent as to its rational status. Thus, FCAT has nothing to say as to whether or not it is rational to form the intention to drink a toxin tomorrow given the offer of a million dollars should one do so before midnight

tonight—see Kavka (1983). When I say that the forming of an intention to perform a given act is not itself an intentional act, I don't mean to be denying either that intentions are attitudes that we form in response to our judgments about reasons or that we can be morally responsible for them. Rather, my claim is only that we don't form intentions at will by intending to form them—see Scanlon (1998, 21-22).

¹³ MT-FCAT is a maximizing view (that is, a view that holds that agents are always required to do what they have most reason to do), for, on MT-FCAT, an act is permissible if and only if the agent has most reason to perform it. It's just that, on MT-FCAT, an agent has most reason to perform an act if and only if it is part of some maximally choice-worthy FCA.

¹⁴ For why we should understand teleological theories as those that make reasons for action a function of reasons to desire rather than value, see Portmore (2009a).

¹⁵ I defend the teleological conception of reasons, and explain how ecumenical it is, in Portmore (2009a and 2009b).

¹⁶ I believe that even FCAs that contain *different* proportions of relaxing acts, altruistic acts, and career-furthering acts can all be maximally choice-worthy. I'll have more to say about this below.

¹⁷ Whereas there is no reason to prefer receiving a cookie today to receiving one on some later day, there is a reason to prefer receiving a fresh-baked cookie today to receiving a pre-packaged one on some later day—or so we'll assume.

¹⁸ If no kind of desire could ever provide an agent with a reason for action, then we should construe 'the preferred option' to stand for 'the option that would provide the agent with the most enjoyment given her present likes and dislikes'. For an argument that the kind of desire that's at play here can provide an agent with a reason for action, see Chang (2004).

¹⁹ This account of rational options is one that I borrow from Bratman (1994, 330-1). See also Portmore (2003, 328-330).

²⁰ I thank an anonymous reviewer for pointing out the need to address this sort of case.

²¹ I borrow the term 'an ultimate end' from Robert Noggle, who uses it to refer to "an intrinsic end that is a fundamental and indispensable part of the agent's life" (2009, 8). An agent's ultimate ends may include such things as the happiness of others, the well-being of her loved ones, and the perfection of herself.

²² The talk of 'enough beneficence' can be misleading. The idea is not that there is some cap on how much beneficence can be required of us such that once we've exceeded this cap we can no longer be required to do anything more for the sake of others. This is implausible, for it seems that no matter how much we have already done for the sake of others, we could still be required to do more, as where, for instance, we happen upon a child drowning in a shallow pond, whom we can save at very little additional cost to ourselves. So although we do need to be sufficiently beneficent to count as having adopted the welfare of others as an ultimate end, the fact that the duty of beneficence consists in the duty to make the welfare of others "a serious, major, continually relevant, life-shaping end" (Hill 2002, 206) means that this duty is never over and done with—see Noggle (2009, esp. 12).

²³ I think that we're also required to adopt other things as ultimate ends—things such as our own perfection and the happiness and perfection of our children.

²⁴ This doesn't mean that that it will be objectively rational for me to continue to put off helping others until my doing so is no longer consistent with my having adopted the welfare of others as an ultimate end. In some instances, it will be in my self-interest to help others sooner rather than later, as where, for instance, I prefer volunteering for Oxfam now and relaxing later to relaxing now and volunteering for Oxfam later.

²⁵ If you have exhibited a pattern of failing to take advantage of good opportunities to help others, then there will come a point at which you can no longer continue to do so and still count as having adopted the welfare of others as an ultimate end. At this point, taking advantage of this opportunity is, given your past neglect, rationally required. Here, then, is another example where the past is relevant on FCAT_i—see note 7.

²⁶ Also, if the example were changed so that I could ensure that Oxfam receives an extra \$200 by sacrificing merely the fulfillment of some very trivial preference (such as, the slight preference I have to write out addresses in print as opposed to cursive), I might, then, be required to make this small sacrifice (thereby sacrificing some very small portion of my self-interest) even if my doing so isn't essential to my leading a full and meaningful life. As I say above, it seems to me that even if one has enough beneficence in one's life to count as leading a life that's full and meaningful, a life that's moderately better in terms of beneficence may outrank a life that's only the slightest bit better in terms of one's self-interest.

²⁷ Again there is no cap on the amount of beneficence that can be required of us. Thus no matter how beneficent one has been in the past, there may be certain "golden opportunities" that everyone is required to take advantage of—see Noggle (2009, 12).

²⁸ I believe that, upon reflection, we will come to think that there are fewer rational options than there initially seemed to be. Also, I admit that many acts that MT-FCAT_i implies are objectively irrational are perfectly rational in the ordinary sense of the word 'rational', for, ordinarily, we reserve the term 'irrational' for only the most egregious actions—see, for instance, Scanlon (1998, 25-30).

²⁹ S performs A_i if and only if S performs each of the act-tokens in that set.

³⁰ Roughly speaking, a reason has justifying strength to the extent that it can make it rationally permissible to perform acts that it would otherwise be objectively irrational to perform, and a reason has requiring strength to the extent that it can make it objectively irrational to refrain from performing acts that it would otherwise be rationally permissible to refrain from performing (Gert 2004).

³¹ On Gert's view, the only reasons that have any requiring strength are those that one has to avoid non-trivial harms to oneself (2004, especially p. 141).

³² Gert says, "It is the official position of this book that altruistic reasons can *never*, in themselves, rationally require action, even though such reasons can justify actions that stand in need of justification" (2004, 89).

³³ For a more precise statement of his view, see Gert (2004, 141). I've added the words 'and because' to Gert's own formulation so as to make clear which side of the bi-conditional has explanatory priority. In correspondence (2/28/09), Gert has indicated that he doesn't object to the inclusion of these words, although I should note that this doesn't represent his final say on the matter.

³⁴ On Raz's view, S is rationally permitted to perform a_i if and only if, and because, S has an undefeated reason to perform a_i . This view is neither substantive nor controversial; indeed, the proponent of FCAT could accept it.

³⁵ Or, at least, this is how he is sometimes interpreted—see Parfit (2008).

³⁶ I thank Robert Noggle for suggesting this case.

³⁷ Equally implausible, to my mind, is that both GT₂ and ST₂ imply that it's rationally permissible for me to spend all of my free time doing nothing but that which would maximize my self-interest.

³⁸ GT₂ is also committed to the view that it would be objectively rational for me to perform an act-series in which I spend all of my free time watching TV—assuming, of course, that this brings no

harm to me. But what fully-informed and well-meaning advisor would recommend that I perform such an act-series given that I would thereby miss out on so many of life's other goods?

Gert clearly thinks that it is rationally permissible to lead an extremely selfish and immoral life—see, for instance, his discussion of the ruthless dictator Carlos Menem (2004, 83), and he argues that intuition is on his side in this matter. I'm not convinced by these arguments, but that's not important. What's important is that intuition is clearly on my side in thinking that it would be objectively irrational to abandon all of one's current projects and relationships so as to dedicate *all* of one's free time either to watching TV or to earning more money to send to Oxfam.

³⁹ This is certainly true if we assume that the only rationally permissible FCAs are those that are maximally choice-worthy, for clearly an FCA in which one inflicts gratuitous self-harm would not be as choice-worthy, other things being equal, as one in which one refrains from doing so. It may seem, though, that we could not account for the irrationality of minor acts of gratuitous self-harm (e.g., gratuitously poking oneself with a needle) if we held that a rationally permissible FCA is any FCA that is good enough, for it may seem that an FCA could contain such an act and still count as being good enough—after all, the harm inflicted is quite minor in proportion to everything else that's included in one's entire future. But this shows only that if we wish to adopt a satisficing version of FCAT, we need to include some restriction to the effect that an FCA counts as good enough only if it doesn't contain an instance of one's performing an act at t_i that makes things worse than they would have been had one instead performed the most nearly effortless act at t_i —the most nearly effortless act being the one that, of all those available to her, requires the least amount of (mental and/or physical) effort. Acts of gratuitous self-harm, even minor ones, would violate such a restriction. And we will also need to include some restriction to the effect that an FCA counts as good enough only if each act contained within it is itself good enough. Acts that entail failing to take easy measures to prevent some harm from befalling oneself would violate this restriction. Of course, once we admit that satisficing FCAT must endorse this restriction, the question becomes: Why not just account for the basic belief by adopting *act-satisficing*—the view that S is rationally permitted to perform a if and only if, and because, it is good enough? The short answer is that act-satisficing would need to be supplemented so as to account for the rational statuses of act-series, and that, as a result, it will face the same sorts of problems that GT and ST encountered.

⁴⁰ The objections that I raise against ST in this section do not concern what ST holds to be rationally permissible. So, unlike the objections that I raised in the previous section, these objections do not necessarily apply to Raz's view. Since Raz never gives us a substantive theory of objective rationality, I'll have to settle for considering ST while admitting that Raz's view may not be subject to all of the same objections.

⁴¹ In an earlier example, I suggested that it would take \$200 to save one or two children. I find this more plausible than Gert's more optimistic estimate. Nevertheless, I'll stick with Gert's numbers when considering his example.

⁴² Gert will count Jane's action as being subjectively irrational, but not objectively irrational. Gert seems to think that it is enough to count this sort of action as subjectively irrational, while allowing that it is objectively rational—see Gert (2007b).

⁴³ It's unclear whether Gert could revise his theory of objective rationality to avoid such counterintuitive implications without in the process undermining one of his main arguments for his thesis that reasons have two separable dimensions of strength. One of the crucial claims that he makes is that it is rationally permissible both to risk one's life and career to save forty unidentifiable children and to refuse to do so—see, for instance, Gert (2004, 88-90; 2003, 10). But if

Gert were to allow that an act's rational status depends on the availability of certain alternatives, then he must deny that it would be rationally permissible to risk one's life and career to save forty unidentifiable children, for ours is a world in which one could save many more unidentifiable children by continuing to live and work while making regular donations to Oxfam. How could it be objectively rational to save only forty unidentifiable children at the cost of one's life when one could instead save many more at the cost of only a few hundred dollars a year?

⁴⁴ See note 8 for clarification.

⁴⁵ I'm assuming that act-utilitarianism is false.

⁴⁶ This is especially problematic for Gert since he holds that an "action is objectively irrational iff virtually everyone would regard the action as irrational," where to regard an action as irrational is to fail to see any features of the action "that could allow someone sincerely to advise someone else to do it" — see Gert (2003, 140). Surely, though, the fact that an act is morally required is a feature of the act that could allow someone to sincerely advise someone else to do it. In his more recent work, though, Gert has suggested that GT might be overly simplistic and that he should, perhaps, count the fact that an act is morally required as a reason that would justify performing it—see Gert (2009).

Regardless, Gert's view (even if revised to allow for the fact that an act's being morally required counts as a reason with justifying strength) implies that it is objectively rational to commit murder and other heinous crimes so long as the benefits that one receives from these crimes compensates for whatever associated risk of punishment there is. The idea that whatever reason one has to refrain from committing murder (and other heinous crimes) has no requiring strength at all is quite implausible.

⁴⁷ This example is borrowed from Portmore (2008).

⁴⁸ As Gert admits, his theory implies that the objective rational status of an act does not depend on what the available alternatives are (2007b, 466-467). Thus if it is objectively rational to incur some harm for sake of saving the life of one child where the only other option is to save no children, then it is objectively rational to incur the same harm for the sake of saving one child even where another option is to save two (or more) children at no additional cost.

⁴⁹ What sort of compromise, if any, it might be rational for me to strike depends on various particulars about me and my situation, such as whether the injuries would prevent me from doing what I most love doing or from fulfilling my other obligations. There may also be more than one reasonable compromise for me to strike, see Subsection 6.4. The problem with ST is that it implies that it would be objectively irrational for any agent to strike any sort of compromise. On ST, each agent is rationally required to save either zero or twelve children.

⁵⁰ Actualism is defended by Goldman (1976), Sobel (1976), Jackson and Pargetter (1986), and Goble (1993). Possibilism is defended by Goldman (1978), Greenspan (1978), Thomason (1981), Humberstone (1983), Feldman (1986 and 1990), Zimmerman (1996 and 2006), and Vessel (2009).

⁵¹ If I believe that I can't control myself and so believe that I couldn't show up and behave myself, then it would be subjectively irrational for me to show up even though I am, on FCAT, objectively rationally required to show up.

⁵² Castañeda (1989) was the first to press this sort of objection against theories such as FCAT. His target was Feldman's theory of absolute moral obligation (Feldman 1986), but the objection applies equally to FCAT, which is modeled on Feldman's theory. For an uncompromising reply to this objection, see Feldman (1990). Currently, I'm content with this sort of uncompromising reply, but it is important to note that should I become discontented with this sort of uncompromising reply, I could revise FCAT so as to better accord with the actualist's intuitions. For how I might do so, see Zimmerman (2006) or Portmore (2009a).

⁵³ The actualist must accept this, for what would happen if I showed up, drank moderately, and behaved myself is better than what would happen if I didn't. See Zimmerman (2006) for more on this point. The example above is a modified version of Zimmerman's example.

⁵⁴ See Portmore (2009a).

⁵⁵ If I'm wrong about this, then we should reject FCAT and accept instead the non-possibilist version of FCAT that I develop in Portmore (2009a), as this theory denies that it's rational to advance a step each week.

⁵⁶ For helpful comments and/or discussions, I thank Ben Bradley, Andrew Cochrane, Peter de Marneffe, Julia Driver, Fred Feldman, Josh Gert, Pat Greenspan, Caspar Hare, Chris Heathwood, Eric Mack, John Maier, G. Shyam Nair, Robert Neal, Robert Noggle, Geoff Sayre-McCord, Dave Shoemaker, Ted Sider, Sergio Tenenbaum, Mark van Roojen, Pekka Väyrynen, Steve Wall, several anonymous reviewers, and audiences at Tulane University, the first annual Rocky Mountain Ethics Congress, and the ninth annual Bellingham Summer Philosophy Conference. For financial support while on a research leave in 2008-09, I thank Arizona State University and the Center for Ethics and Public Affairs, The Murphy Institute, Tulane University.

Works Cited

- Bratman, M. E. 1994. "Kagan on 'The Appeal to Cost'," *Ethics* 104: 325-332.
- Carlson, E. 1995. *Consequentialism Reconsidered*. Dordrecht: Kluwer Publishers.
- Castañeda, H.-N. 1989. "Moral Obligation, Circumstances, and Deontic Foci (A Rejoinder to Fred Feldman)," *Philosophical Studies* 57: 157-174.
- Chang, R. 2004. "Can Desires Provide Reasons for Action?" In *Reason and Value: Themes from the Moral Philosophy of Joseph Raz*, by R. J. Wallace, P. Pettit, S. Scheffler, and M. Smith (eds.), 56-90. Oxford: Clarendon Press.
- Cochrane, A. 2009. "Explaining The Basic Belief." Unpublished manuscript – draft of March 26, 2009.
- Dancy, J. 2004. "Enticing Reasons." In *Reason and Value: Themes from the Moral Philosophy of Joseph Raz*, by R. J. Wallace, P. Pettit, S. Scheffler, and M. Smith (eds.), 91-118. Oxford: Oxford: Clarendon Press.
- Feldman, F. 1990. "A Simpler Solution to the Paradoxes of Deontic Logic," *Philosophical Perspectives* 4: 309-341.
- . 1986. *Doing the Best We Can: An Essay in Informal Deontic Logic*. Dordrecht: D. Reidel Publishing Company.
- Gert, J. 2009. "Moral Reasons and Rational Status." Forthcoming in the *Canadian Journal of Philosophy*.

- . 2008. "Michael Smith and the Rationality of Immoral Action," *The Journal of Ethics* 12: 1-23.
- . 2007a. "Normative Strength and the Balance of Reasons," *Philosophical Review* 116: 533-562.
- . 2007b. "Reply to Tenenbaum," *Canadian Journal of Philosophy* 37: 463-476.
- . 2004. *Brute Rationality: Normativity and Human Action*. Cambridge: Cambridge University Press.
- . 2003. "Requiring and Justifying: Two Dimensions of Normative Strength." *Erkenntnis* 59: 5-36.
- Goble, L. 1993. "The Logic of Obligation, 'Better' and 'Worse'," *Philosophical Studies* 70: 133-163.
- Goldman, H. S. 1978. "Doing the Best One Can." In *Values and Morals*, by A. I. Goldman and J. Kim (eds.), 185-214. Dordrecht: D. Reidel Publishing Company.
- . 1976. "Dated Rightness and Moral Imperfection," *Philosophical Review* 85: 449-487.
- Greenspan, P. 1978. "Oughts and Determinism: A Response to Goldman," *Philosophical Review* 87: 77-83.
- Griffin, J. 1986. *Well-Being: Its Meaning, Measurement, and Moral Importance*. Oxford: Clarendon Press.
- Hill, Jr., T. E. 2002. *Human Welfare and Moral Worth*. Oxford: Oxford University Press.
- Humberstone, I. L. 1983. "The Background of Circumstances," *Pacific Philosophical Quarterly* 64: 19-34.
- Hurka, T. 1990. "Two Kinds of Satisficing," *Philosophical Studies* 59: 107-111.
- Jackson, F. and Pargetter, R. 1986. "Oughts, Options, and Actualism," *Philosophical Review* 95: 233-355.
- Kavka, G. 1983. "The Toxin Puzzle." *Analysis* 43: 33-36.
- Kuhn, S. 2009. "Prisoner's Dilemma." *The Stanford Encyclopedia of Philosophy* (Spring Edition), E. N. Zalta (ed.), URL = <<http://plato.stanford.edu/archives/spr2009/entries/prisoner-dilemma/>>.

- Noggle, R. 2009. "Give Till It Hurts? Beneficence, Imperfect Duties, and a Moderate Response to the Aid Question," *Journal of Social Philosophy* 40: 1-16.
- Parfit, D. 2008. *On What Matters*. Unpublished manuscript – draft of September 28, 2008. Faculty of Philosophy, the University of Oxford.
- Portmore, D. W. 2009a. *Commonsense Consequentialism: Wherein Morality Meets Rationality*. Unpublished manuscript available at <http://www.public.asu.edu/~dportmor/Commonsense%20Consequentialism.htm>.
- . 2009b. "The Teleological Conception of Practical Reasons." Unpublished manuscript available at http://www.public.asu.edu/~dportmor/The_Teleological_Conception_of_Practical_Reasons.pdf.
- . 2008. "Dual-Ranking Act-Consequentialism," *Philosophical Studies* 138: 409-427.
- . 2007. "Welfare, Achievement, and Self-Sacrifice," *Journal of Ethics & Social Philosophy*, www.jesp.org, vol. 2, no. 2.
- . 2000. "Position-Relative Consequentialism, Agent-Centered Options, and Supererogation," *Ethics* 113: 303-332.
- . 2000. "Commonsense Morality and Not Being Required to Maximize the Overall Good," *Philosophical Studies* 100: 193-213.
- Quinn, W. S. 1990. "The Puzzle of the Self-Torturer," *Philosophical Studies* 59: 70-90.
- Raz, J. 1999. *Engaging Reason: On the Theory of Value and Action*. Oxford: Oxford University Press.
- Scanlon, T. M. 1998. *What We Owe to Each Other*. Cambridge, Mass.: Harvard University Press.
- Slote, M. 1989. *Beyond Optimizing: A Study of Rational Choice*. Cambridge, Mass.: Harvard University Press.
- Sobel, J. H. 1976. "Utilitarianism and Past and Future Mistakes," *Noûs* 10: 195-219.
- Tenenbaum, S. 2007. "Brute Requirements: A Critical Notice of J. Gert's *Brute Rationality*." *Canadian Journal of Philosophy* 37: 153-172.

- Thomason, R. H. 1981. "Deontic Logic and the Role of Freedom in Moral Deliberation." In *New studies in deontic logic* by R. Hilpinen (ed.), 177-186. Dordrecht: D. Reidel Publishing Company.
- Velleman, J. D. 1993. "Well-Being and Time." Reprinted in *The Metaphysics of Death* by John Martin Fischer (ed.), 329-3. Stanford: Stanford University Press. The article originally appeared in *Pacific Philosophical Quarterly* 72 (1991): 48-77.
- Vessel, J.-P. 2009. "Defending a Possibilist Insight in Consequentialist Thought," *Philosophical Studies* 142: 183-195.
- Zimmerman, M. J. 2006. "The Relevance of Risk to Wrongdoing." In *The Good, the Right, Life and Death: Essays in Honor of Fred Feldman*, by K. McDaniel, J. R. Raibley, R. Feldman, and M. J. Zimmerman (eds.), 151-170. Burlington, VT: Ashgate Publishing Company.
- . (1996). *The Concept of Moral Obligation*. Cambridge: Cambridge University Press.