

*Love in the Time of Consequentialism**

BARRY MAGUIRE

The University of North Carolina at Chapel Hill

Abstract

There are several powerful motivations for neutral value-based deontic theories such as Act Consequentialism. Traditionally, such theories have had great difficulty accounting for *partiality* towards one's personal relationships and projects. This paper presents a neutral value-based theory that preserves the motivations for Act Consequentialism while vindicating some crucial intuitions about reasons to be partial. There are two central ideas. The first is that when it comes to working out what you ought to do, your friends' interests, the needs of your family, the significance of your own projects and ideals, etc. have *more weight* than the interests and needs of strangers. Your friends' interests are not (thereby) more neutrally valuable than the interests of others. So there is a difference between the value of an outcome and its deontic significance. The second familiar idea is that reasons are modifiable. Reasons of partiality are reasons the weights of which are a function of the value of the relevant outcome modified by facts about the value of caring about the outcome in question. The resulting principle has various further explanatory advantages; in particular, it accounts for project- and relationship-specific permissions and requirements, both at a time and across time.

Introduction

There are several powerful motivations for neutral value-based deontic theories, of which the various forms of Act Consequentialism are the most famous instances. However, Act Consequentialism has great difficulty accounting for a number of pervasive and important ethical phenomena. An important class of phenomena involve *partiality* of one form or another: concerning one's personal relationships and person projects. It is the ambition of this paper to present the structure of an ethical theory that preserves the deepest intuitive motivations for consequentialism while fully vindicating our intuitions about partiality. If Act Consequentialism is defined by its impartiality,¹ I offer a *non-consequentialist* neutral value-based theory of reasons to be partial.

*Special thanks to Gideon Rosen. Thanks to Anthony Appiah, Ralf Bader, Brookes Brown, Anthony Cross, Julia Driver, Antti Kauppinen, Eden Lin, Geoff Sayre-McCord, Matt Kotzen, Kristin Primus, Carlotta Pavese, Chris Melenovsky, Samuel Scheffler, Meghan Flaherty, Carla Merino-Rajme, Karl Schafer, Whitney Schwab, Michael Smith, Mark Johnston, David Velleman, Gerard Vong, Jack Woods, and audiences at The University of Pittsburgh, The Chinese University of Hong Kong, The University of North Carolina at Chapel Hill, The University of St Andrews, The University of Miami, Western Michigan University, Trinity College Dublin, and Princeton University.

I accept weaker principles that capture the core commitments of consequentialism. I accept that facts about the existence of reasons are explained by facts about valuable states of affairs. I also accept that the weight of any reason is a monotonically increasing function of the value of the relevant state of affairs. However, I deny that this value fully explains the weight of the relevant reason.

My solution builds on the familiar idea that when it comes to working out what you have reason to do, your friends' interests, the needs of your family, the significance of your own projects and ideals, etc. have *more weight* than the interests and needs of strangers. Your friends' interests are not (thereby) more neutrally valuable than the interests of others. So there is a difference between the value of an outcome, and its deontic significance.

It is now well known that reasons are modifiable. The weight of a reason can be modified by different kinds of considerations, for instance the extent to which some action of yours would promote some outcome. I will argue, roughly, that reasons of partiality are reasons the weights of which are modified by facts about the value of caring about the outcome in question. On this account, a fact about the value of a state of affairs explains the existence of some reason and something about its weight. A further fact about the value of a distinct state of affairs explains the modification of the weight of the reason. I argue that this value-based account fully vindicates our intuitions about partiality.

Section One: Consequentialism and the Problem of Partiality

The deepest motivation for value-based ethical theories underlies the following famous remark (Singer 1972):

If I am walking past a shallow pond and see a child drowning in it, I ought to wade in and pull the child out. This will mean getting my clothes muddy, but this is insignificant, while the death of the child would presumably be a very bad thing.

The fact that by wading into the pond you can save the life of a child is a reason to do so. The fact that by doing so you'll get your clothes muddy is a reason not to wade into the pond. More generally, if some action available to me would promote some state of affairs that is valuable, the fact that that action would promote that state of affairs is a reason for me to do that thing. This seems to be underlying Peter Singer's response to his example, one of the most well known in ethical theory. In a thesis:

VALUE SUFFICIENCY: If some state of affairs S is valuable, then the fact that some option ϕ would promote S is a reason for you to ϕ .

This sufficiency claim—that facts about value are sufficient for there to be reasons—is a compelling motivation for value-based ethical theories.² Once we have established that some state of affairs is valuable, it is very natural to think that the fact that some action of yours would promote that state of affairs is a reason to perform the action. This is consistent with there being overwhelming reasons not to do it.

I'm talking throughout about *neutral value*. This is a technical name for the familiar idea that value is a monadic gradable property of objects or states of affairs. Hereafter I'll refer to neutral value simply as 'value.'³ I allow the term 'value' to range over instrumental and non-instrumental value.⁴ Some philosophers, faced with the problem of partiality, abandon value for a notion of 'relative value.' They substitute for the monadic property of states of affairs a relation between states of affairs and agents. I take such accounts to abandon the deepest consequentialist insight—exemplified in Singer's example—which is that axiological properties the instantiation of which has nothing to do with the agent, are, as it were, normative for the agent.⁵

I will assume that all reasons can be explained in terms of facts regimented in the form: [ϕ would promote S]. The word 'promote' here is term of art. I will work with the sophisticated account developed in Niko Kolodny (forthcoming), according to which an option *promotes* a state of affairs by instantiating it, causing or partially causing it, constituting or partially constituting it, preventing the preventing of it, or non superfluously probabilifying it. The important detail here is that an option may promote a state of affairs by constituting or instantiating it, for instance when some option constitutes an expression of love, or the perfect pirouette, or an enjoyable afternoon messing around on the river.⁶

Consider now the companion principle:

VALUE NECESSITY: The fact that some option ϕ promotes some state of affairs S is a reason to ϕ only if S is valuable.

This necessity claim is more contentious. One important source of substantive resistance arises from the plausibility of non-value-based 'deontic' reasons, for instance reasons to keep onerous promises to dead relatives, to give parking tickets to the hapless and impoverished, to respect one's friend's voluntary imprudence, to stop at a stop sign in an empty desert. I am optimistic that some version of this principle can be defended against this objection, but for today I'm going to set this class of worries to the side.

I take it that VALUE SUFFICIENCY captures the deepest substantive motivation for value-based views. However, to capture the deepest theoretical motivation, and the purity of the view, one must also defend VALUE NECESSITY. The argument in this essay will help to alleviate some anxiety about the VALUE NECESSITY principle. I will assume, for now, that both are plausible.

These two principles concern the *existence* of reasons. It is natural to add that the *weight* of these reasons is explained, at least in part, by the value of these states of affairs. Since, in Singer's example, the value of the saved life is considerably greater than the disvalue of the muddy clothes, the reason to wade into the pond is considerably weightier than the reason not to. The weight of these reasons is greater, the greater the value of the promoted state of affairs. This thought receives support from a wide spectrum of cases. The more important the charity work, or the better the book, or the more fabulous the festival, the weightier your reasons to participate, and the weightier opposing reasons (for instance the cost involved) would need to be to outweigh these reasons. This also applies to so-called 'deontic'

reasons: plausibly enough, the worse the effect of breaking the law, or the more important it is to tell the truth, the weightier the relevant reason.⁷

When some fact of the form [ϕ promotes S] is a reason, let us call S the *object state*. This line of thought provides motivation for the following principle:

VALUE-WEIGHT CONNECTION: The weight of a reason is a monotonically increasing function of the value of its object state, perhaps along with other variables.

In other words, the weight of a reason is greater, the greater the value of its object state, and lesser, the lesser the value of its object state, other things being equal. This principle is consistent with different monotonically increasing functions from value to weight. It is also consistent with other factors also affecting the weight of a reason in addition to the value of the object state.

So what is the ‘Problem of Partiality’ for value-based theories? The problem begins with an assumption. The plausibility of VALUE SUFFICIENCY and the VALUE-WEIGHT CONNECTION, together with the attractiveness of giving a simple and systematic theory, have encouraged many philosophers to accept VALUE NECESSITY and then to replace the VALUE-WEIGHT CONNECTION with the following stronger principle concerning weight:

IMPARTIAL REASONS: The weight of a reason is fully explained by the value of its object state.⁸

But this package of views—VALUE SUFFICIENCY, VALUE NECESSITY, and IMPARTIAL REASONS—is inconsistent with compelling intuitions about a wide range of cases involving personal relationships, personal projects, voluntary commitments, and ethical ideals.⁹ Such cases seem to provide clear counterexamples to this package of views.

Consider two specific examples. We’ll be working with these throughout:

Relationships: Angela and Becca both love their mothers. Angela’s mother is a beneficent soul, always doing her best to improve the lives of those around her. Becca’s mother is rather self-centered, and never really thinks about or does much for others. Alas, the two mothers are stuck in a burning building. There is only one oxygen mask; only one daughter can save her mother. Becca is holding the mask (she came into its possession by chance). Becca is in a position either to save her mother or enable Angela to save Angela’s slightly more beneficent mother.¹⁰

Projects: “Jack is a lover of art who is visiting the cathedral at Chartres for the first time . . . he is spending much more time at Chartres than he had planned, looking at the cathedral from as many interior and exterior angles, and examining as many of its details, as he can. In fact, he is spending too much time there, from a utilitarian point of view. He had planned to spend only the morning, but he is spending the whole day; and this is going to cause him considerable inconvenience and unpleasantness. He will miss his dinner, do several hours of night driving, which he hates, and have trouble finding a place to sleep” (Adams 1976: 470).¹¹

These cases have the following features. An agent has two available alternatives. One of these alternatives will promote more value or less disvalue than the other.

The alternative that would promote less overall does not promote any particular kind of value in any greater degree than the other alternative. However, the agent is in some sense *partial* to the slightly less valuable action, and not partial to the slightly more valuable alternative. Our question is whether this is a mere impulse, or whether our protagonists have a weightier reason to do the thing they are partial towards or to do the thing that would bring about more value overall.

And here is our datum. In each of these cases the protagonist has a weightier reason to do the thing he or she is partial towards. Becca has more reason to save her mother than to allow Angela to save Angela's mother. Jack has more reason to linger in Chartres cathedral than to go home in time for dinner.

IMPARTIAL REASONS is a natural development of the idea that facts about reasons are to be explained in terms of facts about value.¹² But IMPARTIAL REASONS gets the wrong result in these cases, for the cases are designed in such a way that more value will be promoted by the alternative to the partial action.

Defenders of IMPARTIAL REASONS may attempt to explain, or explain away, all the counterexamples. In some cases, it will indeed be better because of your epistemic situation, or better for everyone in the long run, or better given that you would be more upset than the other daughter, for you to save your mother before saving a stranger. But such replies miss the point. You really have more reason to save your mother! This is not just because you know more about what would make your mother happy or because you are in a better position to make your mother happy. By hypothesis Becca is in the same practical and epistemic position with respect to both mothers in this case.¹³ Intuitively these factors, while perhaps not irrelevant, are not quite to the point. They don't capture all the cases, and in the ones they do capture, they still fail to provide the right explanation of the weight of reasons, and in that sense fail to *vindicate* our partial intuitions. There is a dilemma here. Either the defender of IMPARTIAL REASONS will point to certain non-normative features—knowing one's friend's needs better, being in a better position to help, etc.—which are only coincidentally related to partial reasons and which intuitively come apart from them. Or she will end up having to presuppose what she had hoped to explain, namely that certain partial outcomes are valuable for their own sake. IMPARTIAL REASONS gets the wrong results in our cases, and so we should reject it.¹⁴

It is the goal of this essay to get the right result in these cases. Indeed, more than this, I want to fully *vindicate* the intuitions about reasons expressed in these cases. I want to do so while assuming VALUE SUFFICIENCY and the VALUE-WEIGHT CONNECTION, and also without relinquishing VALUE NECESSITY. Consequently, I will reject IMPARTIAL REASONS and defend a better alternative.

In this paper, I present an original positive proposal. But this is in an important sense a defensive project. The view that any ethical theory starting with neutral value-based principles like VALUE SUFFICIENCY and VALUE NECESSITY will be impartialist is very widely held. I do not know of any alternative account in the literature. My primary aim is to show that there is a range of plausible neutral value-based theories of reasons that can vindicate our intuitions about central cases of partiality. I do not hope to be able to establish that this account can accommodate all of our intuitions about partiality, far less every intuition that

is in tension with Act Consequentialism. Many of our most important ethical convictions are in tension with Act Consequentialism; in particular, I will side-step issues concerning the ethics of promise-keeping, ‘dirty hands,’ and the doctrine of double effect, all of which raise important but slightly different issues better treated separately. This paper is one part of a much larger project, the goal of which is to show that there are resources within the theory of value, reasons, and explanation, to account for a wide range of intuitions ignored by revisionary, toe-the-line consequentialists. I address some of these issues later on, drawing out some further theoretical advantages of my account.

Section Two: Modification in General

Let’s start a brief general discussion of modification with a remark from *The Port-Royal Logic* from 1662:

... to judge what one ought to do to obtain a good or avoid an evil, one must not only consider the good and the evil in itself, but also the probability that it will or will not happen.

This remark can be applied either to ‘objective probability’ or to ‘subjective probability’ or evidence. Suppose you have two drugs, Likely and Unlikely. You know that Likely has an 80% chance of curing the patient, and that Unlikely has a 20% chance of curing the patient, and that everything else is equal. It is hard to deny that you would have more reason to give Likely to the patient than Unlikely. Or suppose that you have two drugs, Reliable and Unreliable. A reliable scientist has told you that Reliable will cure the patient. An unreliable scientist (with a long history of recanting findings) tells you that Unreliable will cure the patient. The consequences of giving neither drug or both to the patient would be an immediate painful death. It is hard to deny that you would have more reason to give Reliable to the patient than Unreliable.

I’m not defending either view here. I’m reminding you that we are already familiar with the business of working out the weights of reasons by means of values and modifiers. The notion of *discounting*, or *attenuating* the weight of some consideration, familiar to students in any introductory microeconomics course, involves just this notion. The account of partiality I’ll offer will rely mainly on intensification—increasing the weight of a consideration—rather than attenuation.

Here’s another example. Prioritarianism is the view that you have more reason to provide the same welfare boost to someone the worse-off they are. Importantly this is not the uncontentious idea that the marginal return to welfare of some good is higher the worse-off the recipient. Rather this is the substantively normative thought that we have more reason to give the equivalent welfare boost to someone the worse-off they are. Again my point for now is not to defend Prioritarianism, but to remind you that the notion of modification, or of a weight-determining function, is familiar and helpful.

There are a few additional reasons why it is helpful for current purposes to think about Prioritarianism.¹⁵ Prioritarianism is explicitly concerned with the weights of

reasons and not with valuable outcomes. Here is how Theron Pummer describes the view: “benefiting people is more morally important, the worse off in absolute terms these people are” (ms: 1). *Being morally important* is a weighted deontic property clearly distinct from the axiological property of being worse off. This principle appears to be inconsistent with IMPARTIAL REASONS, for in a situation in which your options are to give a slightly larger welfare boost to a well-off person, or a slightly smaller welfare boost to a much worse-off person, you may well have more reason to do the latter.¹⁶ Furthermore the modification of weight is explained by an axiological property, namely the absolute welfare of the potential recipient. The full weight of the reason to give some good to some recipient will then be a function of two axiological facts: the welfare boost to the recipient and the prior absolute level of welfare of the recipient.

Section Three: Modification for Partiality

3.1 Modification for partiality

I will argue that the weights of value-based reasons are modified, specifically that they are intensified, by facts about the value of the agent’s partiality towards the outcome.¹⁷ If your being partial towards A is valuable, but your being partial towards B is not valuable (or less valuable), and ϕ promotes A, and φ promotes B, then even if B is more valuable than A, the weight of the reason to ϕ might be greater than the weight of the reason to φ .

Suppose the outcome value of Becca’s mother being saved is 100, and the value of the more beneficent mother being saved is 110. (Of course this is unrealistic precision; it is the structure of the explanation that is important.) Let’s imagine that all other normative considerations are equal, so that there are effectively no other reasons in play. Now consider the value of Becca’s being partial towards her mother being saved and the value of her being partial towards the stranger being saved. I’ll talk about what being partial consists in, and explains the value of partiality, soon enough. For now, let’s say those values are, respectively, 70 and 50. I haven’t said much about the modification function. Let’s assume for now that it simply adds the value of partiality to the value of the outcome. Then we’ll get the modified weights of $100 + 70$ for Becca to save her mother, and $110 + 50$ to save the stranger. Assume that there are no other relevant considerations. Then Becca would have more reason to save her own mother.

The trick is that the partial orientations we are evaluating are the agent’s own. We are not evaluating, for instance, the agent’s partial orientation towards some states of affairs and someone else’s partial orientation towards some state of affairs. It is plausibly of value to be partially oriented towards some outcomes rather than others. This is what makes it possible—even though our explanantia are all neutral values—that agents may have more reason to bring about those states of affairs.

Importantly, I am not appealing to the axiological hypothesis that states of affairs consisting in “mother savings” or “mothers saved by daughters” are neutrally valuable. This would simply yield another reason to throw into the felicific calculation in the standard consequentialist way. In the case of Becca and Angela, if Becca

doesn't save Becca's mother, Angela is poised to save Angela's mother. There will be a mother saving and a daughter saver either way.¹⁸ But our datum is that Becca has more reason to save her mother than to allow Angela to save hers.

3.2 *One substantive account of modification*

This abstract approach is consistent with a variety of different ways of theorising about partial orientations. For the sake of definiteness, I shall offer a particular version. Consider the following thesis, which is the cornerstone of my defense of partiality:

PARTIAL MODIFICATION: The weight of some particular reason of the form [ϕ would promote S] for you to ϕ is greater by some quantity $f(v)$ if and only if, and in virtue of the fact that your caring about S would have value v .

By 'care' I mean roughly what Thomas Hurka means by 'love': "to be positively oriented towards [that] something in one's desires, actions, or feelings, or more generally, in one's attitudes" (2001:13).¹⁹ Here's Agnieszka Jaworska (2007:483) spelling this out a little more:

Typical components of caring include joy and satisfaction when the object of one's care is doing well and advancing and frustration over its misfortunes or setbacks, anger at agents who heedlessly cause such misfortunes or setbacks, pride in the successes for the object and disappointment over its defeats or failures, the desire to help ensure those successes and to help avoid the setbacks, fear when the object is in jeopardy and relief when it escapes untouched, and grief at the loss of the object and the subsequent nostalgia.

What about the *value* of caring? We can also assume, with Hurka, that it is valuable to care about valuable things and disvaluable to care about disvaluable things. It doesn't follow that it is disvaluable not to care about all valuable things. I recognize that there is tremendous value in Medieval tapestries, but I just can't get myself to care about them. I'm missing out, but I care about other things so I don't think this is bad in itself. Our capacity to care is a finite resource, even if many of us don't use it up, and even if the extent of this capacity varies from person to person. Plausibly it is good to care at least a little about everything that is really *important*. It is good to care about every fatal earthquake, however distant. It is also good to be responsive to the value of things we care about: to care much more about the drowning child than the suede shoes. But beyond that, our degree of care for different things will vary significantly. To care *partially* about something is to care about that thing more than you care about other members of its comparison class, correcting for your responsiveness to their respective values.²⁰ (This is compatible with caring about everything considerably more, or considerably less, than your diligent impartial friend.) The value of caring partially is just the value of caring about the relevant thing to the degree that you do relative to other things you care about given your constraints. There are two gradable notions here: the value of caring and the degree of caring. For simplicity in what follows I'll assume that if it is good to care more about S then it is better to care about S.

There are some people such that it would be good if they were impartial, on the whole. Perhaps it would be best if the Pope cared about everyone equally—the next leper just as much as his own mother. But this is not the usual case. As Aristotle said, “Even warm friendship can be felt for only a few” (NE9). For most people most of the time it is good to care more about some people than others.

Consider Becca. She has to decide between saving her mother and enabling Angela to save Angela’s more beneficent mother. It is, intuitively, good that Becca cares about saving her mother more than that she cares about saving the more beneficent stranger.²¹ It is certainly good that she cares about both. But it would be almost monstrous if she were indifferent to whether her mother or a stranger perished in the fire. For each of the two people trapped in the building, there is a perfectly good reason for Becca to save that person from dying in the fire. But it is plausible that Becca has more reason to save her own mother. She also cares more about her mother being saved, and it is good that she does. I am suggesting that the fact that it is good that she cares more about her mother being saved than the stranger explains the fact that she has more reason to save her mother.

Similarly for Jack. Jack is a lover of art. He is greatly excited by the visit to the cathedral, and he is enjoying it enormously:

At the moment, for example, Jack is studying the sixteenth to eighteenth century sculpture on the stone choir screen. He is enjoying this less than other parts of the cathedral, and will not remember it very well. It is not completely unrewarding, but he would have more happiness on balance if he passed by these carvings and saved the time for an earlier departure. Jack knows all this, although it is knowledge to which he is not paying much attention. He brushes it aside and goes on looking at the choir screen because he is more strongly interested in seeing, as nearly as possible, everything in the cathedral than in maximizing utility.²²

Jack cares more about inspecting the screen than about maximizing utility. Moreover, as Robert Adams argues, it would not be better if Jack cared equally and indifferently about anything of value, wherever it might be found. “A great concern to squeeze out the last drop of utility is likely to be a great impediment to the enjoyment of life” (1976:471). He has perfectly good reasons to inspect the screen, and good reasons to leave the cathedral. But he cares more about staying to inspect the screen than about getting home timeously, and it is good that he does. He has more reason to stay and inspect the screen, even though doing so will not maximize utility. This is in line with PARTIAL MODIFICATION.

PARTIAL MODIFICATION states an *explanatory* relationship between the value of caring and the weight of the relevant reason. This is also plausible. It is natural enough to maintain that Jack’s reason to stay in the cathedral, looking at the underwhelming choir screen, is weightier than his reason to go home, precisely in virtue of the fact that it is better that he cares about art than it is that he cares about getting home on time. It is similarly also plausible that Becca’s reason to save her mother is so much weightier in virtue of the fact that her caring about her mother is more valuable.

I am not advocating the subjectivist principle that whatever people *actually* care about modifies their reasons. Such a principle would yield unduly weighty reasons for people to do evil or silly things that they happen to care about, and no particularly weighty reasons to do things in cases where they don't care, but where it would be a great deal better if they did. We could imagine Becca's relationship with her mother being increasingly less a part of what makes her life worth living. Perhaps her mother has been loveless, absent, or abusive. Plausibly, still, Becca has more reason to save her mother than a stranger. But also, plausibly, the weight of her reason to do so diminishes along with the value of caring so much about her mother. Perhaps there would be no value at all in having a special attachment to a loveless and abusive partner. Those who desperately want to support, financially or emotionally, their abusive partners may not have weighty reasons to do so. By contrast those (ingrates) who have no special concern for their loving friends or family may nevertheless have weighty reasons to attend to the interests of those friends or family. Similarly, those who really do just want to take heroin and die, but who could care about something better, and such that it really would be better if they did—their reasons to take heroin are not intensified by the intensity of their passion. Importantly, PARTIAL MODIFICATION maintains that it is the *value* of the attitudes that plays the normative role, not the attitudes themselves.

Section Four: Explaining The Value of Caring Partially

4.1 The value of caring partially

This leads us to the question: What explains the value of caring partially? Consider the substantive thesis that it is better for Becca to care about her mother being saved than a stranger. This is about as close to commonsense as things get in substantive ethical theory. For sure, J.J.C. Smart famously maintained that “the [consequentialist] principle express[es] the attitude of generalized benevolence” (1973: 67). But, as we have seen, it is one thing to say that it is good to care about all creatures, and quite another to say that it would be best to care about each of them equally.

There are many issues here.

Importantly, I am not talking about whether caring about something would be better *for* the agent herself. I am I talking about whether the state of affairs consisting in the agent caring about the object is valuable, and to what degree.²³

Caring partially might be instrumentally valuable or valuable for its own sake. It might be instrumentally valuable for you to care partially about something because you are stuck with it (for instance your lousy committee assignment). It may also be finally valuable to care partially in some cases, e.g. to care about whether your crotchety old parents spend their dotage in comfort and good health. In such cases caring is an essential part of a loving relationship.²⁴

Perhaps the following question is the deepest. Is it plausible that the extrinsic fact about your care for your best friend Julie, that you care about her more than you care about Joe, is valuable for its own sake? Perhaps it is. Certainly, we treasure the ‘specialness’ and privacy and special privileges that we enjoy with our closest

friends and intimate lovers. I'm sure there are restaurants and bridges and even cities you wouldn't want to go to with anybody else. If this extrinsic fact about your care for someone *does* add value to that relationship, then our axiology gives us a clear principled explanation for our reasons to be partial. The explanation of the weight of partial reasons bottoms out, at least in part, in the value for its own sake of caring partially.

It is important to point out that the plausibility of PARTIAL MODIFICATION does not depend on this slightly arcane axiological thesis. There are plenty of other respects in which it will be better to care *more* for some people than for others.

But I suspect that these values are psychologically contingent. If we were all perfect agents, able to interact with any given person equally well, able to engage in any given project equally well, able to perform any given act in any situation without compunction or delay—and if, moreover, we had the resources to engage in any number of relationships, projects, and careful ethical discriminations—if we were all like that, perhaps it would be good if we had impartial attitudes. But clearly we are not such people. We find ourselves with very different predilections, tastes, and talents, and different historical relationships with people: with our family, our nation, the person we sat next to in third grade. We have limited emotional and physical resources, limited time. All of this means that we are bound to form more valuable relationships with certain people than with others, and more valuable commitments to some projects than to others.²⁵

These considerations are similar to the sorts of things that consequentialists mention when they are looking for neutral values in the consequences of partial actions. It is implausible that *every* partial action will promote more value than an alternative. Our cases, Becca's in particular, are specifically designed to the contrary. However, it is plausible that the sorts of considerations appealed to by consequentialists vindicate the greater value of partial *orientations*, evaluated over a period of time. This point is drawn out nicely in the following characterization of the distinctive values of the family, from Harry Brighouse and Adam Swift (2014:85; my italics):

In order to develop into flourishing adults, and to enjoy the goods intrinsic to childhood, children need to have a particular kind of relationship with one or more, but not many more, adults . . . When we say that children need parents . . . we are saying both that there is an essential core to what they need that is best delivered by particular people who interact with them continuously during the course of their development, and that those particular people are able to provide the combination of things needed at any particular time. *Continuity* and *combination* are implied by the idea that that what children need is a particular kind of relationship. It's not the same if different adults take charge of different stages, as if the task of raising children could be passed from one to another like a baton in a relay race. And it's not the same if the different influences that children need are distributed across different adults, so that, for example, the person who loves the child might be someone different from the person who has the discretion paternalistically to substitute her judgement for the child's own. Some change in who does what at different stages, and some functional specialization between different adults at any given stage, may indeed be compatible with children's healthy development. But

there is a particular combination of functions that have to be provided by the same person, or the same few people, for the child's interests reliably to be met.

Parental partiality offers a helpful example of a kind of value that is promoted over time and across a range of activities. Continuity and combination, in different degrees, also enhance the value of your interactions with your mature relatives, your friends, your grocer, your priest, and your dog. The crucial thing about this account of the relevant values is that these values emerge over time and related series of actions. They can be undermined by a single action (e.g. murder), but they cannot be realized by any single action. It is perfectly possible that there will be cases in which, even though you care more about benefitting your child, more value would be promoted by benefitting some other child.

I submit that these claims about the value of caring are plausible, and I'll assume that they are true henceforth. Let me remind you which assumptions I have made so far. I am assuming: VALUE SUFFICIENCY, the VALUE-WEIGHT CONNECTION, that Becca has more reason to save her Mum, and Jack more reason to stay in the cathedral, and now that it is better for Becca to care about her Mum than the stranger, and better for Jack to want to stay in the cathedral than to go home. I am also assuming there is merit in providing an account that is consistent with VALUE NECESSITY. I have suggested that PARTIAL MODIFICATION is consistent with these various assumptions and this ambition, and I have suggested that it may also be a plausible explanation for these data.

4.2 *Actually caring versus being in a position to care*

Consider an objection. Suppose that my father, instead of being a passionate supporter of (i.e. caring a great deal about) the hapless Kilmarnock Football Club, had supported Manchester United, or some other successful club. He would have spent his Saturday evenings celebrating rather than despairing. Suppose that this would be much better overall. Even if so, it is not at all plausible that, right now, he has more reason to go to see Manchester United than Kilmarnock FC. He has no interest in Manchester United! This feels like a counterexample: it would be better to care more about Manchester United than Kilmarnock, but he has more reason to see Kilmarnock than Manchester United. Isn't that inconsistent with PARTIAL MODIFICATION?

In response, I reject the claim that it would be better if my Dad, here and now, cared more about Manchester United. Maybe it would *have* been better if he *had* cared about Manchester United, and it would be better if he *were to* do so in the future, but it wouldn't be better here and now.

The simplest argument for this would be that he could care about Manchester United right now only by being another person altogether. But this is clearly too strong. There are some things such that it would be good if you cared about those things, and such that you really wouldn't be a different person if you did care about that thing: e.g. whether Idris Elba is cast as James Bond.

However, there is a weaker sense in which you would 'be a different person' if you didn't care about—to change the example—horses but instead cared about

cats. Many facts about you affect what it is good that you care about right now. You grew up with horses, you have spent hundreds of hours grooming and riding and talking about horses. Since you have a history of emotional and physical investment in horses, and not cats, it is better for you to care about horses than cats right now. It might be better for you not to care so much about horses, given that you went bankrupt, lost your stables, and are utterly heartbroken. A lifetime with cats would have better prepared you for this unforeseeable disaster. But even though, now, it might be better for you to *come to* care about cats rather than horses, it is not true of you, here and now, in the face of this disaster, that it would be better if you cared about cats. For you know nothing about cats. You don't even like cats. Perhaps this will change, and perhaps you should work on changing it. But, right now, it is best for you to care about horses. Similarly, even when they are losing, it is better that my father cares about Kilmarnock rather than Manchester United. He has a history with Kilmarnock, of loving and losing (and very occasionally winning). For sure, if things had been different, he might have lived a happier life as a Manchester United fan. Maybe he should even work, over time, to change his loyalties. But right now things are the way they are, and he really does care about Kilmarnock. And who knows, maybe they'll win next week.

This line of thought extends to people who don't care about things it would be good if they cared about. For it is various facts external to the care itself that make caring about one thing better than another. It is the fact that you have groomed horses, ridden horses, that your family ride horses, that you have a big group of friends that you meet up with at horse shows, all of whom are mad about horses. These facts do not suffice to make it the case that it would be good if you cared about horses. Perhaps you are indoorsy, and bookish, and more into reptiles. But for many people, in that situation, it would be natural, and indeed good, for them to care about horses.

4.3 *The directness of caring*

I can introduce another advantage of the appeal to care at this point, concerning the *directness* of caring about someone or something. Caring about something does not essentially involve believing anything about whether that thing is valuable, or whether one's caring about it is valuable, or whether it is valuable that good things are happening to it. Rather when you care about something you put yourself emotionally and motivationally 'at its disposal.' As Jaworska nicely puts it, you see the thing as "commanding [your] emotional vulnerability" (2007:484). People who really care about each other become "subject to one another" (*Ephesians*, 5:21). There isn't a good word for this in English. In the polyamorous community, the word "compersion" is sometimes used to describe a state like this. Perhaps this is also what Octavio Paz had in mind when he said in *Piedra De Sol* that "to love is to undress our names."

Caring fits us together.²⁶ Token caring attitudes manifest a particular way of relating directly to other things that really matter to us—our future selves, our families, friends, careers, our passions. Clearly it is good, in some sense, to be directly emotionally sensitive to significant events involving people and things we

care about—to feel joy when they feel joy, and even to suffer when they suffer. It has even been suggested—alongside Jaworska, by many philosophers including Bernard Williams and Harry Frankfurt—that what you care about plays a large part in determining who you are.

This is precisely the kind of thing we need to explain partiality. Partiality, after all, essentially involves a fact about the agent who is partial – recall that to care partially about something is, in part, to care about that thing more than you care about other members of its comparison class—and not just about the object of partiality. We need a fact about the agent that affects the weights of reasons but without self-indulgence or ‘one thought too many.’ You can act on the fact that you care about someone without anything like self-indulgence or ‘one thought too many.’²⁷ You can do something for someone because you care about *that person*. Facts about you are part of the full explanation of your motivation, but this is not self-indulgence. This explanation has precisely the profile we would want from an explanans.

These different factors—the values of the outcomes, and the significance of an agent’s attitudes—intuitively play different roles in the determination of what that agent has most reason to do. The kind of explanation I propose is ‘non-horizontal.’ The value of Becca’s partiality towards her mother is not just another evaluatively significant consideration to be summed with the rest and weighed against other neutral values among the outcomes, as IMPARTIAL REASONS would have it. As Bernard Williams points out, “of course [the agent’s] feelings, so regarded, are of virtually no weight at all in relation to other things at stake” (1973:103). The problem here is with the theory not with the data. The problem is with this way of regarding the agent’s feelings. According to PARTIAL MODIFICATION the value of her partiality plays a different role: it intensifies other reasons she has.

Now to the question: *why* does the weight of your reason to do something depend on the value of your caring about the outcome? The first pass answer is: the fact that you care about the outcome plausibly affects the weight of your reason. The facts about what we care about are deeply important facts about ourselves and our relationships with each other. This is a feature with the right profile, depth, and familiarity to play this key role in ethical explanation. But it is well known from objections to subjectivisms of various kinds that actual psychology fares poorly as a non-revisionary explanation of our reasons. What we want instead is facts about caring that meet some axiological standard; we want facts about the *value* of caring. Caring is also an inter-temporal state, that ranges over different situations. Allowing such a state, meeting some axiological standard, to play the modifying role provides continuity in our reasons for action over time.

Section Five: Partiality and Deontology

5.1 *Set aside constraints and dirty hands*

As noted at the outset, the primary goal of this paper is to respond to a specific challenge to value-based ethical theories, namely that they cannot account for the weight of reasons for actions in the context of personal relationships and projects. This is one of the most compelling objections to Act Consequentialism. But Act

Consequentialism is only one specific instance of a family of value-first ethical theories, and there are resources beyond those appealed to by Act Consequentialism that can accommodate and even vindicate partialist intuitions.

I have focused on providing an account of the weights of reasons. I have said nothing so far about *obligations* to loved ones, special *duties* to one's friends and one's family, moral or other *requirements*, *prohibitions*, and *permissions*, or whether certain actions are *optional* or *supererogatory*. A fully satisfactory ethical theory must give an account of these properties.

Not all of these properties are related to partiality. There are various *constraints* on the performance of certain kinds of actions (to use McNaughton and Rawling's term (1998)), for instance the breaking of promises, or lying, or violating property rights, or interfering with rights to freedom of speech. I am prohibited from interfering with your expression of political views that are objectionable by my lights.²⁸ These constraints apply independently of my personal projects and specific relationships, and they consequently fall beyond the scope of the current essay. There are also morally significant differences pertaining to degrees of *agential involvement* in bringing about outcomes (to use Wedgwood's (2011) helpful term). I may have more reason against harming someone than failing to benefit that person, or against killing someone than against letting them die. Since these cases raise different sets of issues, I will also set them aside.²⁹

5.2 Distinguishing reasons from what they explain

One version of the 'Partialist Intuition' maintains that it is permissible to do something in promotion of a project or the interest of a loved one, rather than performing some value maximizing alternative. This is also an important motivation for the hypothesis that certain actions are supererogatory, where this claim is construed as the thesis that it is sometimes not the case that one ought to perform one's value-maximizing option. Hence one strong motivation for idea that certain actions are supererogatory is accommodated by my account. I also deny that one ought always to maximize value, precisely in cases in which one is partial (in the relevant sense) to some alternative (and the values at stake are thus and so).³⁰

My strategy seems to account for the permissibility of failing to take the option that maximizes value. Still, one might reasonably complain that there isn't really a *permission* to do something other than maximize value in cases like Angela and Jack. You have most reason to take the partial option rather than to maximize value in those cases. There is not yet any leeway in the account I propose. It's a bit like the pirate saying you have permission to walk the plank.

One response to this is to point out that by rejecting the correlation between impartial value maximization and what one ought to do, I have made maximization a little less unattractive. It is more plausible that one ought to do what one has most reason to do if the weights of one's reasons are sensitive in the ways I suggest to what it is good to care about.

I think that is right as far as it goes, but I will make an alternative response. Here I will emphasize another theoretical advantage of working with a Value-Based Theory of Reasons rather than Consequentialism. Distinguish two projects.

One is that of giving an account of the existence and weight of reasons. It is the goal of this paper to provide accounts that analyze these notions in terms of facts about neutral value and other non-normative facts (such as facts about what it is to care, and facts about promotion). A second project is to give an account of the relationship between facts about reasons (and their weights) and various strict deontic facts such as the all things considered ought, and also the notions of permissibility, requirement, prohibitions, obligation, criticizability, etc.³¹ This second project will account for optionality, maximization, the supererogatory, etc. on the other hand. The account of the weights of reasons on offer here is compatible with a range of views about the relationships between reasons and other deontic facts. To give an example, the debate between maximizing theories and satisficing theories falls under this second project.³² Strictly speaking, I'm agnostic about this second project, and hence about maximization, optionality, etc. For all I've said, it might be permissible for you to fail to do what you have most reason to do. In this paper I have made no commitments so far about the relations between facts about reasons and their weights, on the one hand, and facts about ought, permissibility, requirements, prohibitions, or obligations,³³ on the other.

However, this response also only goes so far. Although some of these deontic facts are explanatorily downstream from facts about reasons, others are not. Clearly, there is a range of facts about permissions, prohibitions, and requirements that themselves explain the weights of our reasons in the context of personal relationships and projects. By and large, many of the imperatives and permissions and prohibitions in life are internal to projects or relationships of one kind or another. You *cannot* miss your niece's birthday party; you *may* skip the Connemara Horse Show; you *must* go to the Dublin show; my Dad *must* see the quarter-final; but he *could* miss the play-off, since he lost the bet he is *obligated* to bring me along. Such properties as these are central to our normative lives but often ignored by normative theories in the spirit of Act Consequentialism or IMPARTIAL REASONS. In the remainder of this section I will introduce two further features of PARTIAL MODIFICATION that enable it to do better than Act Consequentialism in accounting for these phenomena.

5.3 *Internal deontic properties specific to options*

Here is the first feature. To participate in a relationship or a project is, in part, to accept a host of relationship- or project-specific permissions, prohibitions, and requirements. These relationship and project-internal requirements, permissions, and prohibitions affect the value of caring about specific outcomes (as distinguished from the individuals or projects more generally). Consequently, via PARTIAL MODIFICATION, these deontological facts affect the weights of your reasons.

Consider an overly simple deontological view, according to which the following relations hold between relationship- and project-based permissions, requirements, etc. and the all things considered ought: that if you have a relationship-specific permission to ϕ then it is not the case that you ought not ϕ , and that if you are required to ϕ then you ought to ϕ . Then if you have a relationship-based permission to ϕ , it is not the case that you ought not ϕ . But this definitely seems too strong. There are two points to make here. One is that even perfectly significant internal

considerations might be outweighed by other more significant considerations. You may have a relationship-dependent permission to tell me a vulgar joke, but if the madman will blow up Manhattan if you do so, clearly you ought not do so. This is a familiar point, even if not uncontentious. (This point also applies to constraints: it is permissible for you to steal my bike to chase down the madman.) These permissions, prohibitions, and requirements, while perfectly genuine and significant, are not *all things considered* permissions, prohibitions, and requirements, that are explained, *inter alia*, by the weights of reasons. They are themselves *contributory* considerations. They will often be weighed against each other, or against non-partial normative considerations, such as the saving of children from ponds. The second point is that evil projects and destructive relationships and projects also yield internal deontic properties: if you are collecting scalps, you will be required to have a clean, sharp blade, you will be prohibited from being overly sentimental about your victims, and perhaps you will be permitted to discard the bald ones. You will, of course, be required to collect a reasonable number of scalps on every expedition. But it does not follow that you ought to do so.

So we should think that project- and relationship-specific deontic properties have a *derivative* and *contributory* normative status. The normative significance of some such relationship-specific requirement or permission is partly explained by the normative significance of the relationship itself. In turn, we should think that the normative significance of some action that constitutes participation in a project or relationship is a function of *both* of these standards of assessment: the internal assessment (in this case, the strength of the requirement or permission) and the external assessment of the project or relationship itself.

The view I propose attends separately to these different levels of explanation. The value of caring about some specific state of affairs is a function of the value of the more general care, and the *significance* of this state of affairs to this more general care. Given that you care about comic books, it would be good to care a great deal about catching this bus, since it is the only way to get to the convention. Given the nature of your relationship with your son, it is good that you care a great deal about him studying for school—and perfectly fine not to care about the finer details of his sex life. The criterion for significance will vary from relationship to relationship, project to project. Sometimes this will be a matter of straightforwardly promoting the relevant values. Often it will be a matter of adhering to relationship- or project-specific commitments.

We can use the theoretical flexibility provided by these different levels of analysis to generate plausible results in a range of cases. It is clearly better, in general, to care about your teenage daughter than your record collection. But imagine a situation in which your daughter is suffering some minor inconvenience and your entire collection is at risk of being destroyed. It is plausibly good that you care more avoiding the destruction of your record collection than about failing to help your daughter with her homework. Since I'm assuming that if it is good to care more about S than R then it is better to care about S than R, it follows that it is better to care about the imminent destruction of your record collection than about the difficulties in your daughter's homework.³⁴ By PARTIAL MODIFICATION, you may

well have more reason to attend to your record collection than your daughter. Abstractly: you may have more reason at some time to make a more significant contribution to a less valuable project or relationship.

5.4 *Internal deontic properties over time*

Act Consequentialism and IMPARTIAL REASONS both explain deontic facts at a time just in terms of the values promoted by the various alternatives available at that time. With Bernard Williams and many other critics of Act Consequentialism, I find this time-slice centrism objectionable.³⁵

By contrast, PARTIAL MODIFICATION appeals to *generalized attitudes*, which are evaluated over a period of time. This appeal to diachronic facts imposes some stability upon your reasons over time by increasing the weight of your reasons to do things to it would be good if you were positively oriented towards over time, and de-emphasizing your reasons to do other things. This spreads the source of the explanation for what you have most reason to do at a time across other times, to your relationships, commitments, and valuable projects and valuable dispositions. This provides a kind of continuity in your reasons over time.

The significance of some specific state of affairs to some general care may also be sensitive to something like your standing with respect to the outcome. This passage from Elizabeth Anderson offers a helpful illustration of another kind of permissiveness built into many of our relationships (1993:151):

The norms of gift exchange differ from the norms of market exchange in several respects . . . Although both forms of exchange involve reciprocity, the form and timing of the return of good differ in the two cases. In market exchange, an uncontracted delay in reciprocation is cause for legal action. But the exchange of gifts among friends usually incorporates an informal understanding of reciprocity only in the long term. To be anxious to “settle accounts” of small sums, as when one person insists upon splitting a restaurant tab exactly in half, calculating sums to the penny, is to reject the logic of friendship.

It is not good to be concerned to reciprocate a friend’s kindness immediately. Indeed, it is good to be unconcerned about this, since that bespeaks a more robust acceptance of these more permissive reciprocal norms.³⁶

There are two kinds of permissibility in play now. The permissibility of failing to care about your son’s sex life is *option-specific*. The permission pertains to the option in question in virtue of facts about that option together with the facts about the nature of your relationship with your son. By contrast, the permissibility of reciprocating a favour tomorrow rather than today is *diachronic*. The latter kind of permissiveness has the same structure as ‘imperfect duties,’ the distinguishing feature of which is that they do not confer obligations or requirements in any particular situation.³⁷ Instead they obligate one to perform *enough* of a certain kind of action, or to contribute *enough* to some end, over some reasonable period of time. Anderson points out, by contrast with reciprocation in the marketplace, that the reciprocation of favours in a friendship has this imperfect structure. With many projects and relationships, it is good that you care about contributing *enough* to any

given project over time. It doesn't follow that you need to care about contributing as much as you can at every opportunity. To do that is to reject the 'logic of friendship.' Caring about art is also like this. You have some time off this afternoon and you'll be in the center of London. Should you spend the time in the National Gallery? But you've been visiting museums all week, and you sat through *Götterdämmerung* last night. It might be perfectly fine to be relatively unmoved by the prospect of several hours wandering around the National Gallery this afternoon, fantastic as it is, and even if spending those hours in the National Gallery would promote more value than spending the afternoon walking by the Thames.³⁸

This is the converse of the earlier point about Jack spending longer in Chartres Cathedral. If he were at the end of a long week of sightseeing, it might well be best that he cares most about getting home early and getting a good night's sleep. Caring about art doesn't mean caring about every single opportunity to experience art, since you care about other things as well. Sometimes you should care more about things that you should care about less. The paradoxical air disappears when we distinguish caring about something in general and caring about some particular instance of it.

I have given an account of the weight of reasons to be partial. I have suggested that these weights can be affected by two kinds of relationship- or project-specific permissions, requirements, and prohibitions. One kind is specific to option-types, the other is sensitive to one's standing over time. It is worth explicitly noting some limitations of the strategy so far developed. I have said nothing about what happens when different projects conflict. You might have cause to care about your daughter qua daughter, qua friend, qua scrabble opponent, qua defendant. For instance, when one is a judge, the fact that the defendant is one's daughter does *not* seem to be grounds for partial treatment. The question of how to these different modifiers interact must be left to another day.

Section Six: Two Objections

6.1 *Objection one: modifying reasons or grounding reasons?*

At this point, we can consider another challenge. Suppose you are convinced that valuable caring about some outcome is normatively relevant, and more specifically you are convinced that it increases the total weight of reason in favour of doing the thing that promotes that outcome, and that it does so in a way that accommodates a range of deontological intuitions. Still there are various ways that the value of caring about an outcome might affect the weight of reason to promote that outcome, and you might prefer some alternative to PARTIAL MODIFICATION.

I have been arguing that valuable caring intensifies extant reasons, that is, reasons the existence of which is explained by something else, namely the value of the outcome. But why not maintain that valuable caring about an outcome is a sufficient condition for the *existence* of reasons? Why not think that valuable caring about an outcome generates *sui generis* reasons, rather than modifying other reasons? In other words, why not accept the following principle:

CARE SUFFICIENCY: If your caring about the obtaining of some state of affairs S would be valuable, then, in virtue of that fact, the fact that some action ϕ would promote S is

a reason for you to ϕ ; the weight of this reason is a monotonically increasing function of the value of caring about S.

This is an indirect principle, in the sense in which Rule Consequentialism is an indirect principle. Since it is not a necessity principle, it is consistent with the provision of direct reasons as well. I confess that I find this principle attractive. This view, just like PARTIAL MODIFICATION, can, in principle, yield the right results in our working cases. Moreover, for CARE SUFFICIENCY, just as for PARTIAL MODIFICATION, reasons and their weights are explained in terms of axiological facts and other non-deontic facts (apart from the function itself).

In the end, I won't be too disappointed if you accept CARE SUFFICIENCY instead of PARTIAL MODIFICATION. The conclusion of this paper is that we can explain reasons of partiality with neutral values, and either principle can do this. It is not clear that much hinges on this difference. However, as I shall now explain, some broadly theoretical considerations incline me to prefer, and to recommend, the package consisting of VALUE NECESSITY, VALUE SUFFICIENCY and PARTIAL MODIFICATION to the package of CARE SUFFICIENCY alongside VALUE SUFFICIENCY.

6.2 *Caring about valueless outcomes*

CARE SUFFICIENCY may be inconsistent with VALUE NECESSITY. VALUE NECESSITY maintains that all reasons are explained by facts about valuable outcomes. CARE SUFFICIENCY maintains that some reasons are explained not by facts about valuable outcomes but rather by facts about the value of caring about an outcome. There is an explanatory disagreement here, which I will address shortly. There would also be an extensional disagreement only if the following principle were false: that it is valuable to care about a state of affairs only if that state of affairs is valuable. It is plausible that this principle is false.³⁹ Moreover, as I'll now argue, whenever it is good to care about a valueless outcome, one's reasons do not correlate with the value of caring about the outcome, but with the value of the outcome, and hence in such cases have no weight. This is the first consideration against CARE SUFFICIENCY.

I have been assuming that when one cares about x , one wants good things for x —one wants improvements in the well-being of people one loves, progress towards the realization of one's goals and ideals, etc. I also noted that it might not be neutrally valuable that these things are promoted. There are bad goals, bad ideals, and bad people. It is plausibly good that the mother of the diabolically evil son cares whether he suffers in jail, even if it is good that he suffers. But I don't think she has reasons to keep him out of jail—and that's because her reasons are tracking the values of the outcomes, not the value of her caring about the outcomes.

Or take Jack's case. It is not quite true that no value whatsoever will be promoted by studying the choir screen, even when he is tired. But such a variation on the case wouldn't be too far-fetched. I'm sure every art lover has wasted half an hour in a predictably dreadful art gallery at some point. I find it plausible that, in such cases, the existence of reasons is counterfactually sensitive to the facts about the value of the outcomes, not to the facts about the value of caring. If there really would be *nothing* good about staying and inspecting the choir screen, then Jack plausibly wouldn't have any reason not to head home.

Perhaps it could be argued that if there really were nothing good left to see (as Jack was well aware) then it wouldn't be good that Jack wanted to stay. Although it would be *understandable* for him to be disappointed by leaving valueless parts of the choir screen unexamined, it really isn't good that that would disappoint him. Recall that the value of caring about some particular outcome is a function both of some fact about the value of caring generally about the thing (an ideal, in this case), together with some fact about the *significance* of this outcome in that general context. When Jack knows that there is nothing good left to see, his reluctance to leave might be understandable, but it is not itself valuable. This line of thought invites the stronger suggestion that it is valuable to care about a state of affairs only if that state of affairs is valuable.

In some cases, *extrinsic* value will be realized by a seemingly valueless outcome. There may well be value promoted by hanging around in boring parties with friends, or reading mediocre poetry, or drinking barely passable home-brew. To return to Brighthouse and Swift's example, participating in distinctively parental activity of the sort described partly constitutes the realization of particular kinds of 'family values.' Plausibly, participation in such activities as these has constitutive value of some kind. It is plausible that even reading boring bedtime stories realizes constitutive part value; this is less likely in the case of traipsing around watching bad performance art. You have no reason to look at completely worthless art, whether you are a lover of art or not.

To sum up: it is plausible that caring about x involves being positively oriented towards good things for x . This doesn't entail that whenever it is good to care about x being F , it would be good if x were F . This entailment doesn't go through even though there are various ways in which it might be good if x were F , e.g. extrinsically good. Either way, it is plausible that the existence of reasons to promote the relevant state of affairs track along with the neutral value of the outcomes. Hence accepting VALUE NECESSITY rather than CARE SUFFICIENCY is more theoretically flexible. It is also theoretically simpler—since we can insist that all reasons are explained by neutrally valuable outcomes. It is also theoretically more conservative. The existence every reason is explained directly by the promotion of valuable states of affairs. This is closer to the letter and the spirit of IMPARTIAL REASONS and Act Consequentialism. Only the weights of reasons are explained indirectly. As we have seen, Act Consequentialists and others are not ideologically opposed to various kinds of modification. These considerations provide support for the package of VALUE NECESSITY, VALUE SUFFICIENCY and PARTIAL MODIFICATION rather than the package of VALUE SUFFICIENCY and CARE SUFFICIENCY.

There is another reason for wanting to hang on to VALUE NECESSITY. The position I have developed has some similarities with indirect consequentialisms such as Rule Consequentialism. There are differences, of course; for instance, I do not claim that *all* reasons are explained indirectly. Plenty of reasons are explained directly by the values of outcomes, for instance the two reasons in Singer's example (the drowning child and the muddied suede). Indeed, one attractive feature of the present account is that it undermines the standard dichotomy between direct and indirect versions of consequentialism. Still, insofar as my account has some

affinities with indirect consequentialism, it is worthwhile pointing out that it avoids the standard objection to such views, namely the ‘rule worship’ objection.⁴⁰ According to that objection, Rule Consequentialism advocates actions even in cases where it is obvious that the very thing following rules is designed to promote—the maximization of value—will be undermined by doing so. By hanging on to VALUE NECESSITY we retain as a necessary condition for the existence of reasons that some value would be promoted by performing the action in question.⁴¹

6.3 *Objection two: is this view really value-first?*

There is something objectionably *flattening* about Act Consequentialism, and about IMPARTIAL REASONS. Every significant consideration gets thrown into the felicific calculation, and you ought to do the thing that gets the highest score, irrespective of anything else about you—your values, your commitments, your relationships, etc.—except insofar as these are evaluatively significant parts of the consequences. Otherwise your deontic significance, as one might put it, is that you happen “to be near certain causal levers at a certain time” (Williams 1981: 4).

The alternative I have proposed—reasons explained by the value of outcomes, and weights intensified by the value of orientations towards them—has more structure. This makes room for different kinds of considerations to play different deontic roles. I have argued that the considerations playing the higher-order role are neither primitively deontic facts nor merely psychological facts, but rather facts about the value of orientations, specifically the value of caring about outcomes.

I have argued that modification is familiar and generally considered to be an acceptable part of a value-first approach. My account is a little more complex than either Direct Consequentialism or Indirect (Attitude) Consequentialism, since my account maintains that the weights of reasons are explained by the explanans of both: partly by the value of outcomes and partly by the value of attitudes. This complexity itself is not problematic, but it raises one further ideological commitment. For according to the view I have proposed—the package of VALUE NECESSITY, VALUE SUFFICIENCY, and PARTIAL MODIFICATION, the weights of reasons are explained by facts about the values of outcomes and the values of attitudes *together with some monotonically increasing function* that assigns weights to these considerations. Different theorists may advocate different functions, and these will have direct and dramatic substantive implications. More impartially inclined theorists will give partial attitudes less significance; more partially inclined theorists will give partial attitudes greater significance.

I take this to be a theoretical strength of the approach. These alternatives should not be ruled out *analytically*, but rather substantively. Does it follow that this package does not provide a strictly *value-based* account of weight, since it appeals in addition to this weighting function? I’m not sure. It is possible that the ‘balancing’ of valuable attitudes and valuable outcomes is underwritten by some higher-order value, perhaps a value instantiated by actions, such as the value of fairness or justice. But at this point the notion of ‘value’ seems to have become less distinctive; this would be a theoretical cost.

Instead, I am inclined to accept the conclusion of the objection, that my account involves some primitive normative non-axiological principle. Perhaps some Act Consequentialisms that reject deontic facts altogether—Alastair Norcross (2006) comes to mind—avoid any non-axiological normative commitments. But still, I think I'm in good company. Most Act Consequentialisms involve some substantive principle relating facts about values with deontic facts. The distinction between maximizing and satisficing Act Consequentialisms, for instance, turns on facts about the function from value to ought, where these facts are not determined by facts about value.

Conclusion

Traditional value-based deontic theories have relied on a proportionality thesis: that the weights of reasons to do something are fully explained just by the neutral value promoted by that thing. A range of cases involving partiality of one kind or another plausibly refutes this proportionality thesis. I have proposed that the weight of some reason is a monotonically increasing function of two values: the value of the outcome and the value of the agent's orientation towards the outcome. I argue more specifically that there are advantages to understanding orientation in terms of caring about the outcome. This approach is more complex than Act Consequentialism. But anyone who is not an Act Consequentialist usually finds the very simplicity of Act Consequentialism to be a mark against its plausibility, not a mark in its favour. I motivated the dual sensitivity of the weight of reasons to these two kinds of value. I also pointed out some further advantages of this approach, in providing diachronic stability in our reasons over time, accommodating various deontological data about permissions and requirements at a time and over time, and vindicating the directness of our responsiveness to those we care about. A final motivation for this approach is that it is theoretically ecumenical, drawing insights from theories often held to be in competition. It preserves the main motivations for consequentialism, namely VALUE SUFFICIENCY, VALUE NECESSITY, and the VALUE-WEIGHT CONNECTION. It systematically accommodates important motivations for deontological views, arising from the deontic significance of special relationships, personal projects, and voluntary commitments. It also draws on some insights from the ethics of care literature and virtue ethics, by giving special (though not overly special) normative significance to valuable attitudes, and to our deepest and most valuable cares and concerns. Many questions of detail and substance remain, and I have not had space to provide responses to alternative views here. But I hope to have shown that there is plenty of theoretical room for value-first ethical theorists to defend plausible principles that vindicate the partialist motivations characterized at the outset, and which are interesting and fruitful in their own right.

Notes

¹For this taxonomical claim, see Krister Bykvist (2002), Michael Smith (2003), and Ralph Wedgwood (2009).

²On the idea that this principle captures the intuitive core of consequentialism, see W. Frankena (1963), *Ethics*, chapter two.

³ I'll assume that having more or less value or disvalue correlates with being better or worse. I'll often refer to the disjunction, 'value or disvalue' with the single term 'value'—the context will disambiguate. For readability, I'll often use 'good,' 'bad' and cognates as synonyms for 'valuable,' 'disvaluable,' etc.

⁴ For some helpful distinctions, see Ralf Bader's "Two Value-Theoretic Trichotomies" (ms).

⁵ On the idea that consequentialism is distinctively concerned with *neutral* values, see Frank Jackson, 1997. On the notion of *relative value* see Michael Smith (2003) and Douglas Portmore (2001, 2003).

⁶ It falls to the theory of value, not the theory of promotion, to determine which actions are valuable for their own sake. In some cases, this will seem ad hoc, e.g. if one attempted to argue that there is more value in keeping one's own promise than enabling others to keep their own promises. The account developed here is sympathetic to this worry about "consequentializing" deontologically significant properties. On the notion of consequentializing, see Dreier (1993), Smith (2009).

⁷ On this gambit, see Michael Philips (1987) and Walter Sinnott-Armstrong (2009).

⁸ More plausibly: the weight of a reason is a monotonically increasing function in two variables: the value of its object state and the extent to which the action would promote the object state.

⁹ I address the scope of the phenomenon under consideration in Section Five. The challenge until then is to give a plausible vindication of central intuitions in cases involving relationships and projects.

¹⁰ This is a modified version of the famous Fenelon and the valet case in W. Godwin (1793).

¹¹ Intuitions about this case may vary. Perhaps some other case would have made the point slightly more sharply. I stick with this case out of deference to Robert Adams, since the insights in his 1976 paper inspired many of the ideas in this paper. I use lots of other examples as we go through, e.g. my Dad's support for Kilmarnock F.C.

¹² This thesis is *direct*: the deontic status of an action is directly proportionate to the evaluative status of its outcome. In this respect this thesis is similar to *Act Consequentialism*, which I take to be the thesis that one ought to do whatever would maximize neutral value. However, there are two important differences between IMPARTIAL REASONS and Act Consequentialism. One is that IMPARTIAL REASONS is concerned one by one with evaluatively valenced states of affairs, each of which constitutes a separate reason to perform the action in question. Secondly IMPARTIAL REASONS is not a maximizing thesis. It takes no stand on the question of whether you ought to do what you have most reason to do. Consequently, it is consistent with accounts of ought, accounts which make room for supererogation, satisficing, etc. I discuss this further in Section Five.

¹³ In a well-known defense, Frank Jackson argues that "the special regard that we have for a relatively very small group of people—to the extent that it is morally justified—[can] be explained probabilistically, in terms of our special epistemological status with regard to our nearest and dearest" (1991:473; compare Sidgwick, 1907:431). This maneuver will work for some cases, but not all. The Angela and Jack cases are designed to avoid any relevant informational asymmetries. Compare Susan Wolf (1992:244).

¹⁴ Since Global Consequentialism (cf. Peter Railton 1984) entails Act Consequentialism, it also gets the wrong answers about what one ought to do in these cases.

¹⁵ Cf. Parfit (1997), Hooker (1994, 2010), Pummer (ms).

¹⁶ I say 'appears,' since this Prioritarian principle is Anti-Consequentialist unless it could somehow be argued that, in each case, giving extra weight to someone the worse off they are promotes some value distinct from welfare, for instance the value of *promoting equality* (an action) (cf. Parfit *op.cit.*). If this were true, perhaps this could be modelled differently: You have reasons to promote the welfare of each, and an additional reason to promote equality or fairness, the value of which is proportional to the increase in the weight of your reason to provide the benefit to the individual how is absolutely worse-off. I'll address an analogous issue applicable to my own account shortly.

¹⁷ For an earlier attempt, see Maguire (2012). For a brief but helpful discussion of a similar gambit, see Ralf Bader (2016). Thanks to Ralf for helpful discussions.

¹⁸ Compare Simon Keller (2013:106).

¹⁹ 'Love' would have better suited my title, but I worry that it would have an overly narrow connotation. 'Care' may already have a slightly narrower connotation than I need.

²⁰ This analysis involves three ideas: some account of the appropriate comparison class; some specific *comparative* psychological fact, the fact that you care about one thing *more* than others; and the correction for their *value* as members of the kind in question, assuming your attitudes are responsive to

these values. Here's further support for the third condition. Suppose that one member of a comparison class is clearly better than the rest: Giulietta Masina in a line-up of street mimes. You like each member of the class precisely in proportion to their value as mimes, and consequently you like Giulietta Masina the most. I don't think we would say you are *partial* to Giulietta Masina. Suppose you like Giulietta Masina the most, but you actually like her a little less than her value as a mime merits, while you like one of the others more than her value as a mime merits. I don't think we would say you are partial to Giulietta Masina, because you like her the most, but rather that you are partial to this other one mime, because you like her more than her value as a mime merits.

²¹ Notice that it is the value of caring about the *outcome* that plays the relevant deontic role, not the value of caring about the action. Usually, insofar as you care about someone, you care about her getting some good, not exactly about *bringing about* the outcome in question. Sometimes the action will *be* the valuable outcome, for instance when you are spending time sitting at the nursing home with your grandmother. In such cases then the value of caring about the action is deontically significant because the action is the relevant outcome.

²² Adams, *op.cit.* p. 471.

²³ Throughout, I restrict myself to the slightly less idiomatic 'it is good that A cares about *x*' instead of 'it is good for A to care about *x*' to emphasize that this is a point about neutral value not welfare.

²⁴ An interesting question remains about the 'wrong kind' of instrumental value: e.g. if an evil demon provides a sufficiently robust (over time and counterfactuals) incentive for you to care about *x*. That would make it valuable for you to care about *x*, but that fact intuitively wouldn't affect the weights of your reasons to promote good things for *x*. An example of one strategy: assume an analogue of 'wrong kinds of reasons scepticism': that the evil demon really makes it valuable to *appear to the demon* as though you care about *x*.

²⁵ Here is a slightly different argument for the value of caring partially. Think again about the Frank Jackson's (*op. cit.*) defense of the view that, *in every relevant case*, more expected value will be promoted by performing the action one is partial to, for epistemic and situational reasons. This is implausible. But it is not at all implausible that this will be true about the value of *caring* for someone, where this is an enduring state evaluated over a significant period of time and range of counterfactual circumstances.

²⁶ I elaborate on the relationship between caring and fittingness in my (ms).

²⁷ Cf. Williams (1973) and Markovits (2010). Issues surrounding the 'one thought too many' objection are too complex to be treated here in detail. In brief, I will make four points. (1) A theory of reasons rather than ought is already in a position to account for the normative significance of this or that fact rather than some single overall fact about value maximization. (Contrast with Markovits 2010: 231). Moreover, according to the value-based theory of reasons developed in Maguire (2016), facts about value are not reasons but parts of the grounds of reasons facts. A similar strategy is pursued in Schroeder (2007). (2) According to PARTIAL MODIFICATION, facts about caring (meeting some axiological condition) partly explain why you ought to act partially. This captures some of the Williamsian data. But in any case, (3) the account developed here is a metaphysics of reasons, not a theory of reasoning. Nothing follows from this account about how one should think about what to do. I reckon that PARTIAL MODIFICATION does a better job at bringing together the criterion of rightness and the right decision making procedure than 'Global Consequentialism.' But the fact remains that more work needs to be done to develop any challenge to my account here. (4) Williams surely has one thought too few. In one's quieter moments it is ethically responsible to reflect on the importance of what one cares about, and to encourage oneself to concentrate more on the things that are more important, and less on the less. In non-emergency moral predicaments, even those involving loved ones, it is plausible that responsible moral agents should think about moral theory in the course of their careful moral deliberation. Thanks to Keshav Singh for helpful discussion.

²⁸ For a value-theoretic account of rights, See Philip Pettit (1988).

²⁹ For criticism of the deontic significance of agential involvement, see Scheffler (1982). For an interesting positive proposal, see Wedgwood (2011).

³⁰ In fact, I account for this motivation without taking a stand on the question of how best to analyze the property of being supererogatory, and without taking a stand on the question of whether such actions are supererogatory. For a seminal discussion, see Urmson (1958). For a recent discussion, see Portmore (ms).

³¹ For further discussion of this two level strategy, see Maguire (2016). For important recent discussions of this second level, specifically in the context of weighing moral reasons, see Harman (forthcoming) and Schafer (2016). For alternatives to this ‘two project’ approach see Gert’s distinction between ‘justifying’ and ‘requiring’ reasons (2016) and Raz’s account of ‘exclusionary reasons’ (1990).

³² On satisficing, see Pettit and Slote (1984). For criticism, Bradley (2006).

³³ Samuel Scheffler (2010) has argued that there is a deontically significant difference between partiality in relationship and partiality in projects, consisting in the fact that relationships but not projects generate obligations. This is not a distinction in the weight of associated reasons. It is more like a distinction in the nature and conditions of fitting criticism for non-compliance. Consequently, this distinction falls beyond the scope of the PARTIAL MODIFICATION principle, which is concerned just with weights of reasons. Thanks to Samuel Scheffler for discussion.

³⁴ Thanks to an anonymous referee for encouraging me to think through cases with this structure.

³⁵ For an analogous concern in single person cases, see Warren Quinn (1990) and Tenenbaum and Raffman (2012).

³⁶ The logic of friendship is of course not always permissive. You are required to get your sister a birthday present, to take care of your sick child, to attend your best friend’s wedding.

³⁷ Except the last one: if you have an imperfect duty to see an opera a month and its the last day of the month and you haven’t seen one yet, then you have a perfect duty to do so. Thanks to Gideon Rosen for discussion.

³⁸ For a similar approach focusing on first-person cases, see Tenenbaum and Raffman, *op. cit.*

³⁹ *Pace* Thomas Hurka *op. cit.* I assumed a version of this principle earlier for the sake of argument. I am now weakening that assumption.

⁴⁰ Cf. Smart (1973).

⁴¹ Is this enough? Don’t we need a reason of significant weight? Well, the weight of the reason is also sensitive to the value of the outcome. That should be enough.

Works Cited

- Adams, R. 1976. ‘Motive Utilitarianism.’ *The Journal of Philosophy* 73(14): 467–481.
- Anderson, E. 1993. *Value in Ethics and Economics*. Harvard University Press.
- Arnauld, A. 1964. *The Art of Thinking, Port-Royal Logic*, transl. Dickoff & James. Library of Liberal Arts.
- Bader, R. 2016. ‘Conditions, Modifiers, and Holism’ in *Weighing Reasons*, (eds) Lord & Maguire. Oxford University Press.
- . (ms). ‘Two Value-Theoretic Trichotomies.’ New York University.
- Bradley, B. 2006. ‘Against Satisficing Consequentialism.’ *Utilitas* 18(2): 97–108.
- Brighouse, H. & Swift, A. 2014. *Family Values*. Princeton University Press.
- Bykvist, K. 2002. ‘Alternative Actions and the Spirit of Consequentialism.’ *Philosophical Studies* 107(1): 45–68.
- Cowan, T. 2007. ‘Caring About the Distant Future.’ *The University of Chicago Law Review* 74(1): 5–40.
- Dreier, J. 1993. ‘Structures of Normative Theories.’ *The Monist* 76(10): 22–40.
- Frankena, W. 1963. *Ethics*. Prentice-Hall.
- Gert, J. 2016. ‘The Distinction Between Justifying and Requiring: Nothing to Fear’ in *Weighing Reasons*, (eds) Lord & Maguire. Oxford University Press.
- Godwin, W. 1976. (1793). *Enquiry Concerning Political Justice*. Penguin.
- Harman, E. (Forthcoming.) ‘Morally Permissible Moral Mistakes.’ *Ethics*.
- Hooker, Brad, 1994. ‘Is Rule-Consequentialism a Rubber Duck?’ *Analysis* 54(2): 92–97.
- . 2010. ‘When is Impartiality Morally Appropriate?’ in *Partiality and Impartiality: Morality, Special Relationships, and the Wider World*, (eds) Feltham & Cottingham. Oxford University Press.
- Hurka, T. 2001. *Virtue, Vice, & Value*. Oxford University Press.
- Jackson, F. 1991. ‘Decision-Theoretic Consequentialism and the Nearest and Dearest Objection.’ *Ethics* 101(3): 461–482.
- Jackson, J. 1997. ‘Which Effects’ in *Reading Parfit*, (ed) Dancy. Oxford University Press.

- Jaworska, A. 2007. 'Caring and Full Moral Standing.' *Ethics* 117(3): 460–497.
- Keller, S. 2013. *Partiality*. Princeton University Press.
- Kolodny, N. (forthcoming). 'Instrumental Reasons' in *Oxford Handbook of Reasons and Rationality*, (ed) Star. Oxford University Press.
- McNaughton, D. & Rawling, P. 'On Defending Deontology.' *Ratio* 11(1): 37–54.
- Maguire, B. 2012. 'Values, Reasons, & Ought.' Phd Dissertation, Princeton University.
- . 2016. 'The Value-Based Theory of Reasons.' *Ergo* 3(9).
- . (ms). 'There are No Reasons for Emotions or Desires.' The University of North Carolina at Chapel Hill.
- Markovits, J. 2010. 'Acting for the Right Reason.' *The Philosophical Review* 119(2): 201–242.
- Parfit, D. 1997. 'Equality and Priority.' *Ratio* 10(3): 202–221.
- Paz, O. 'Piedra de Sol' or 'Sunstone.' Translated by Eliot Weinberger and Muriel Rukeyser.
- Pettit, P. & Slote, M. 1984. 'Satisficing Consequentialism.' *Proceedings of the Aristotelian Society*, Supplementary Volumes 58: 139–163, 165–176.
- . 1988. 'The Consequentialist Can Recognise Rights.' *Philosophical Quarterly* 38(150): 42–55.
- . 1989. 'Consequentialism and Respect for Persons.' *Ethics* 100(1): 116–126.
- Philips, M. 1987. 'Weighing Moral Reasons.' *Mind* 96(383): 367–375.
- Portmore, D. 2001. 'Can an Act-Consequentialist Theory Be Agent Relative?' *American Philosophical Quarterly* 38(4): 363–77.
- . 2003. 'Position-Relative Consequentialism, Agent-Centered Options, and Supererogation.' *Ethics* 113(2): 303–332.
- . (ms). 'Transitivity, Moral Latitude, and Supererogation.' Arizona State University.
- Pummer, T. (ms). *The Priority Monster*. The University of St. Andrews.
- Quinn, W. 1990. 'The Puzzle of the Self-Torturer.' *Philosophical Studies* 59(1): 79–90.
- Railton, P. 1984. 'Alienation, Consequentialism, and the Demands of Morality.' *Philosophy and Public Affairs*. 13(2): 134–171.
- Raz, J., 1990. *Practical Reason and Norms*. Princeton University Press.
- Scheffler, S. 1982. *The Rejection of Consequentialism*. Oxford University Press.
- . 2010. 'Morality and Reasonable Partiality' in *Equality and Tradition*. Oxford University Press.
- Sidgwick, H. 1907. *The Methods of Ethics*. McMillan & Co.
- Schafer, K. 2016. 'The Modesty of the Moral Point of View' in *Weighing Reasons*, (eds) Lord & Maguire. Oxford University Press.
- Schroeder, M. 2007. *Slaves of the Passions*. Oxford University Press.
- Singer, P. 1972. 'Famine, Affluence, and Morality.' *Philosophy & Public Affairs* 1(3): 229–243.
- Sinnott-Armstrong, W. 2009. 'How Strong is This Obligation.' *Analysis* 69(3): 438–442.
- Smart, J.J.C. 1973. *Utilitarianism: For and Against*. Cambridge University Press.
- Smith, M. 2003. 'Neutral and Relative Value After Moore.' *Ethics* 113(3): 576–598.
- . 2009. 'Consequentialism and the Nearest and Dearest Objection' in *Minds, Ethics, and Conditionals: Themes from the Philosophy of Frank Jackson*, (ed) Ravenscroft. Oxford University Press.
- Tenenbaum, S. & Raffman, D. 2012. 'Vague Projects and the Puzzle of the Self-Torturer.' *Ethics* 100(1): 116–126.
- Urmson, J. 1958. 'Saints and Heroes' in *Essays in Moral Philosophy*, (ed) Melden. University of Washington Press.
- Wedgwood, R. 2009. 'Intrinsic Value and Reasons for Action' in *Metaethics*, (eds) Sosa & Villaneuva. Wiley Periodicals, Inc.
- . 2011. 'Defending Double Effect.' *Ratio*, Special issue on Deontological Ethics (ed. Hooker): 384–401.
- Williams, B. 1981. 'Persons, Character, and Morality' in *Moral Luck*, (ed) Rachels. Cambridge University Press.
- Wolf, S. 1992. 'Morality and Partiality.' *Philosophical Perspectives* 6: 243–259.